

3060



US006480850B1

(12) **United States Patent**
Veldhuisen

(10) Patent No.: **US 6,480,850 B1**
(45) Date of Patent: **Nov. 12, 2002**

(54) **SYSTEM AND METHOD FOR MANAGING DATA PRIVACY IN A DATABASE MANAGEMENT SYSTEM INCLUDING A DEPENDENTLY CONNECTED PRIVACY DATA MART**

(75) Inventor: **Adriaan W. Veldhuisen**, San Marcos, CA (US)

(73) Assignee: **NCR Corporation**, Dayton, OH (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **09/411,337**

(22) Filed: **Oct. 1, 1999**

Related U.S. Application Data

(63) Continuation-in-part of application No. 09/165,784, filed on Oct. 2, 1998, now Pat. No. 6,253,203.

(51) Int. Cl.⁷ **G06F 17/30**

(52) U.S. Cl. **707/9; 707/2; 707/201**

(58) Field of Search **707/1, 2-5, 7-9, 707/10, 100-102, 500, 526-527, 200-206; 705/35-37, 38; 713/200-202; 709/217-219**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,319,777 A • 6/1994 Perez 707/1

5,751,949 A • 5/1998 Thomson et al. 713/201
5,961,593 A • 10/1999 Gabber et al. 709/219
6,085,191 A • 7/2000 Fisher et al. 707/9
6,141,658 A • 10/2000 Mehr et al. 707/7
6,195,657 B1 • 2/2001 Rucker et al. 707/5
6,253,203 B1 • 6/2001 O'Flaherty et al. 707/9
6,275,824 B1 • 8/2001 O'Flaherty et al. 707/9

* cited by examiner

Primary Examiner—Alford W. Kindred

(74) Attorney, Agent, or Firm—James M. Stover

(57)

ABSTRACT

A system for managing data privacy comprises a database management system for storing data from a plurality of consumer database tables, with irrevocable logging of all access, whether granted or denied, to the data contents stored in the consumer data tables; a privacy metadata system that administers and records all data, users and usage of data that is registered as containing privacy elements; and a replication system that feeds the consumer access system with personal consumer data, maintains integrity of the consumer data and provides changes and corrections back to the originating database management system through their own integrity filters as well as a means of storage and the mechanism to provide input for changes in the personal data or privacy preferences. The system further includes means for managing consumer notification, access, correction and change of preferences for privacy or data protection in the privacy metadata system.

6 Claims, 11 Drawing Sheets

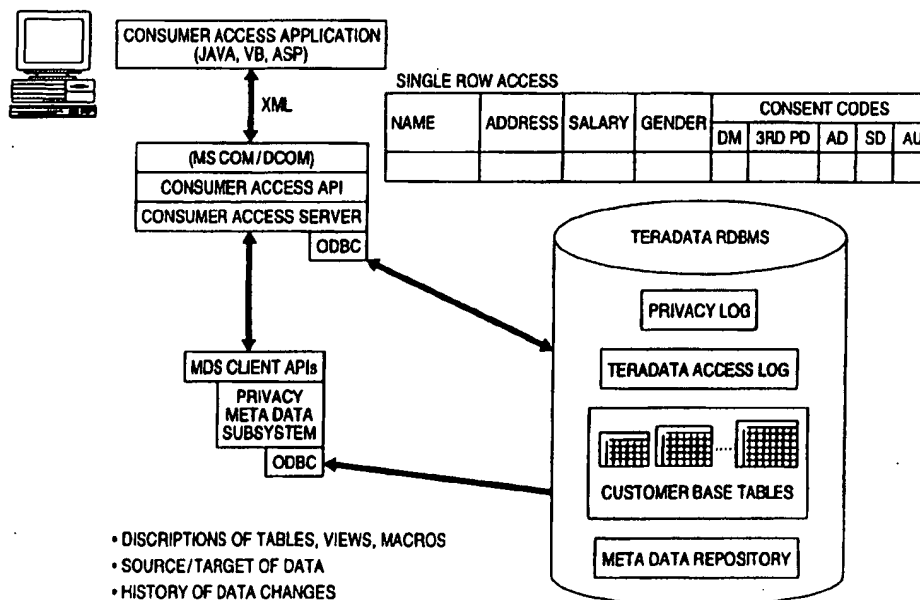


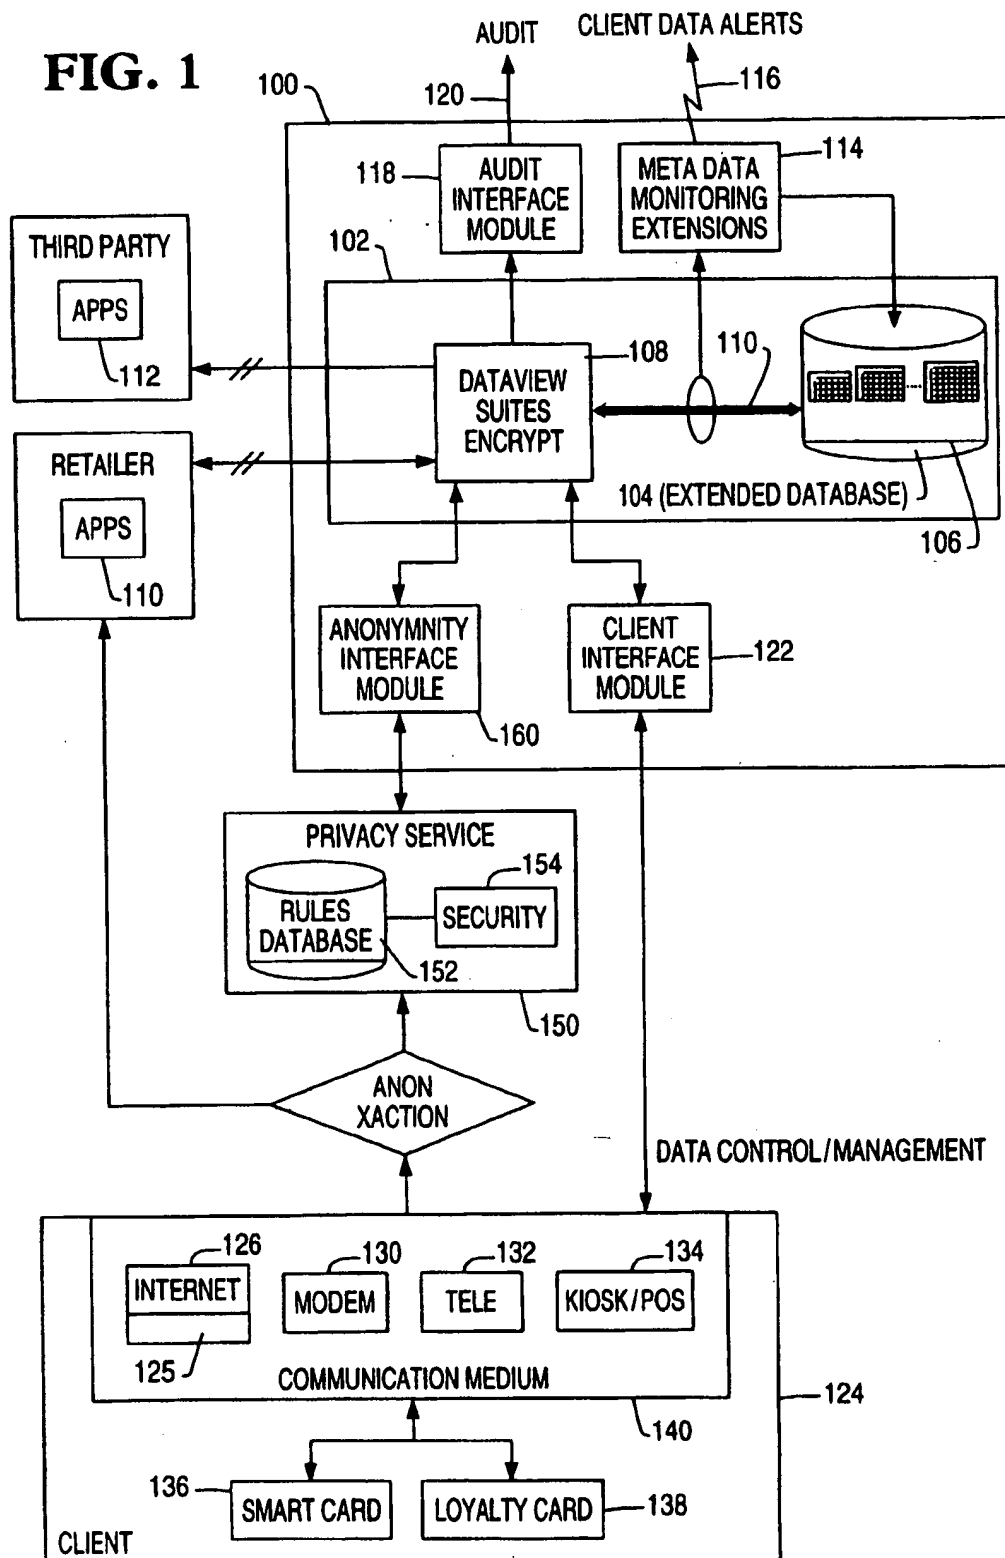
FIG. 1

FIG. 2A

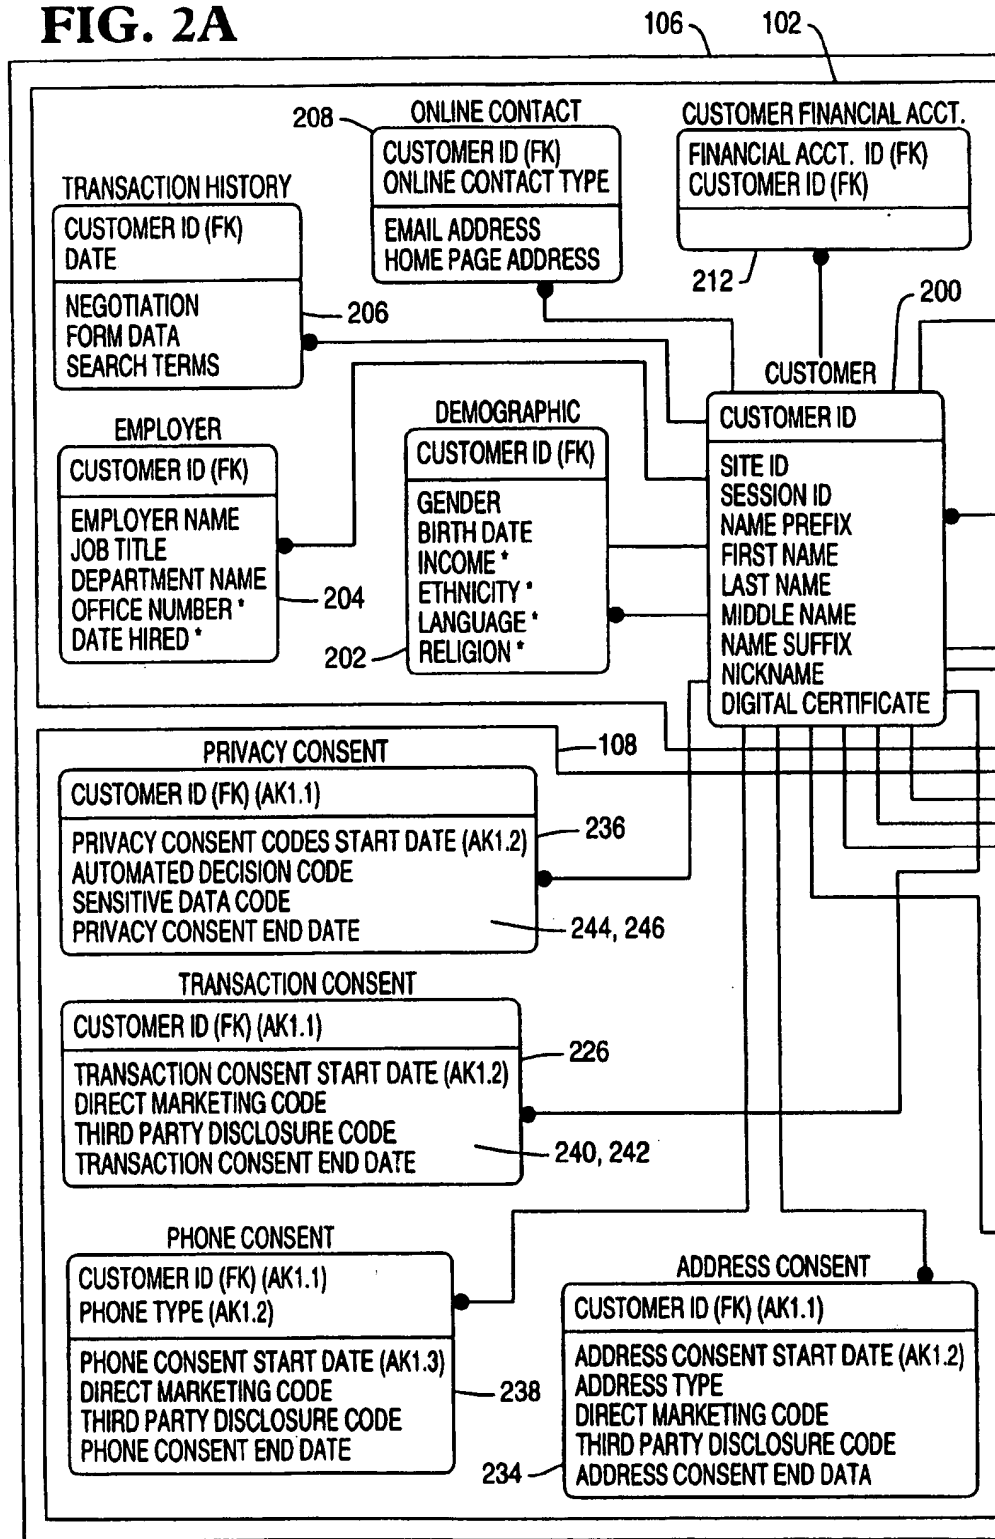
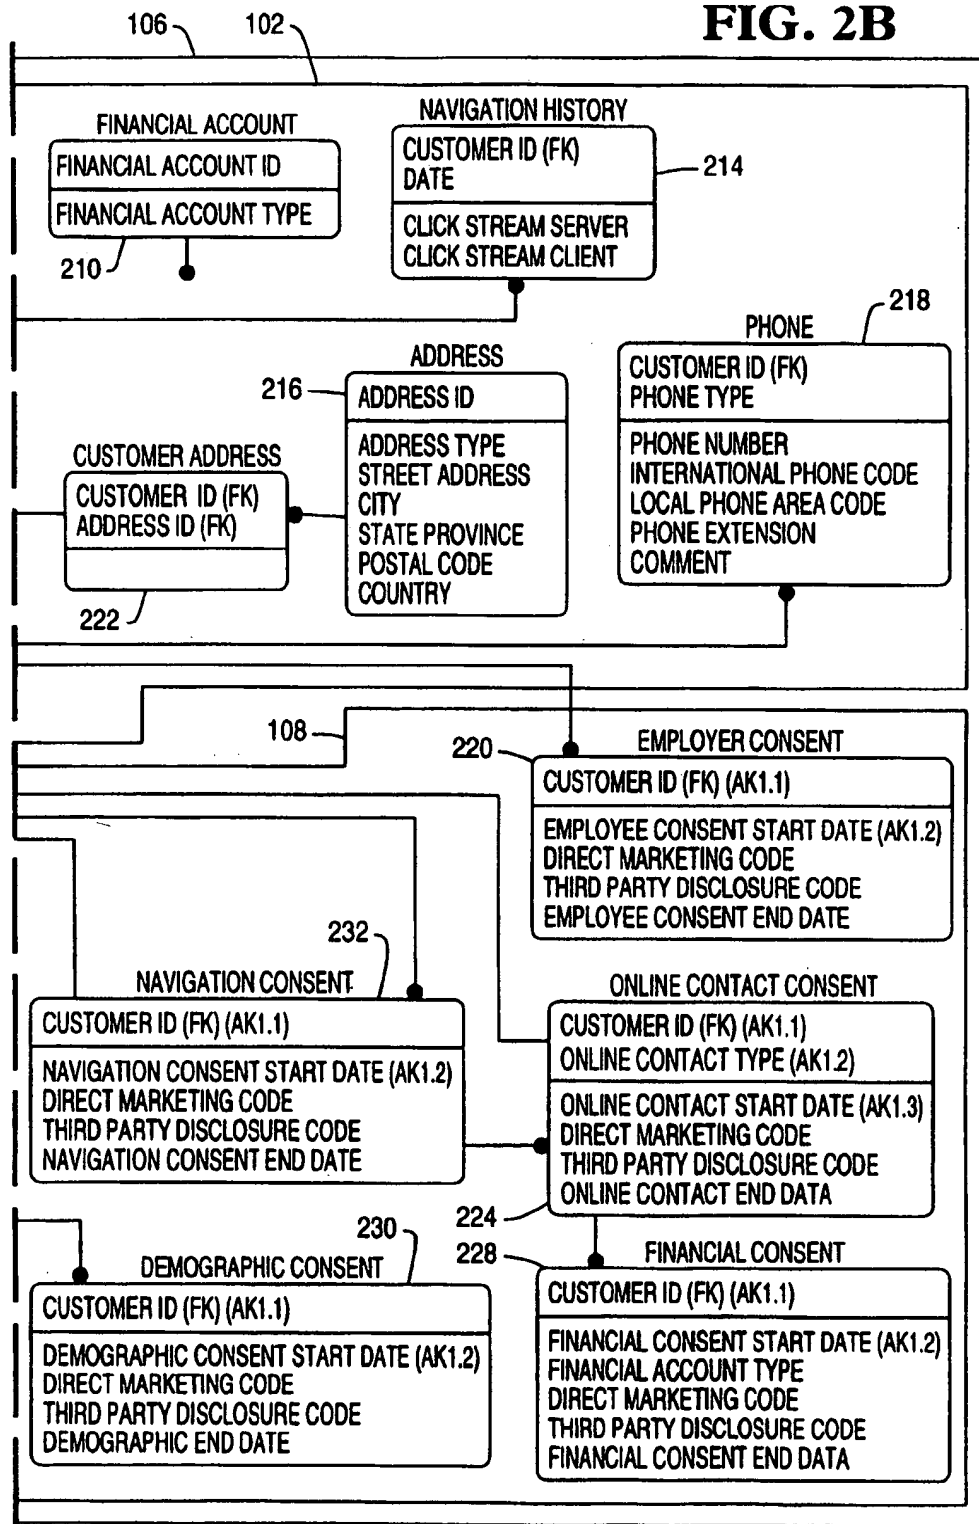


FIG. 2B



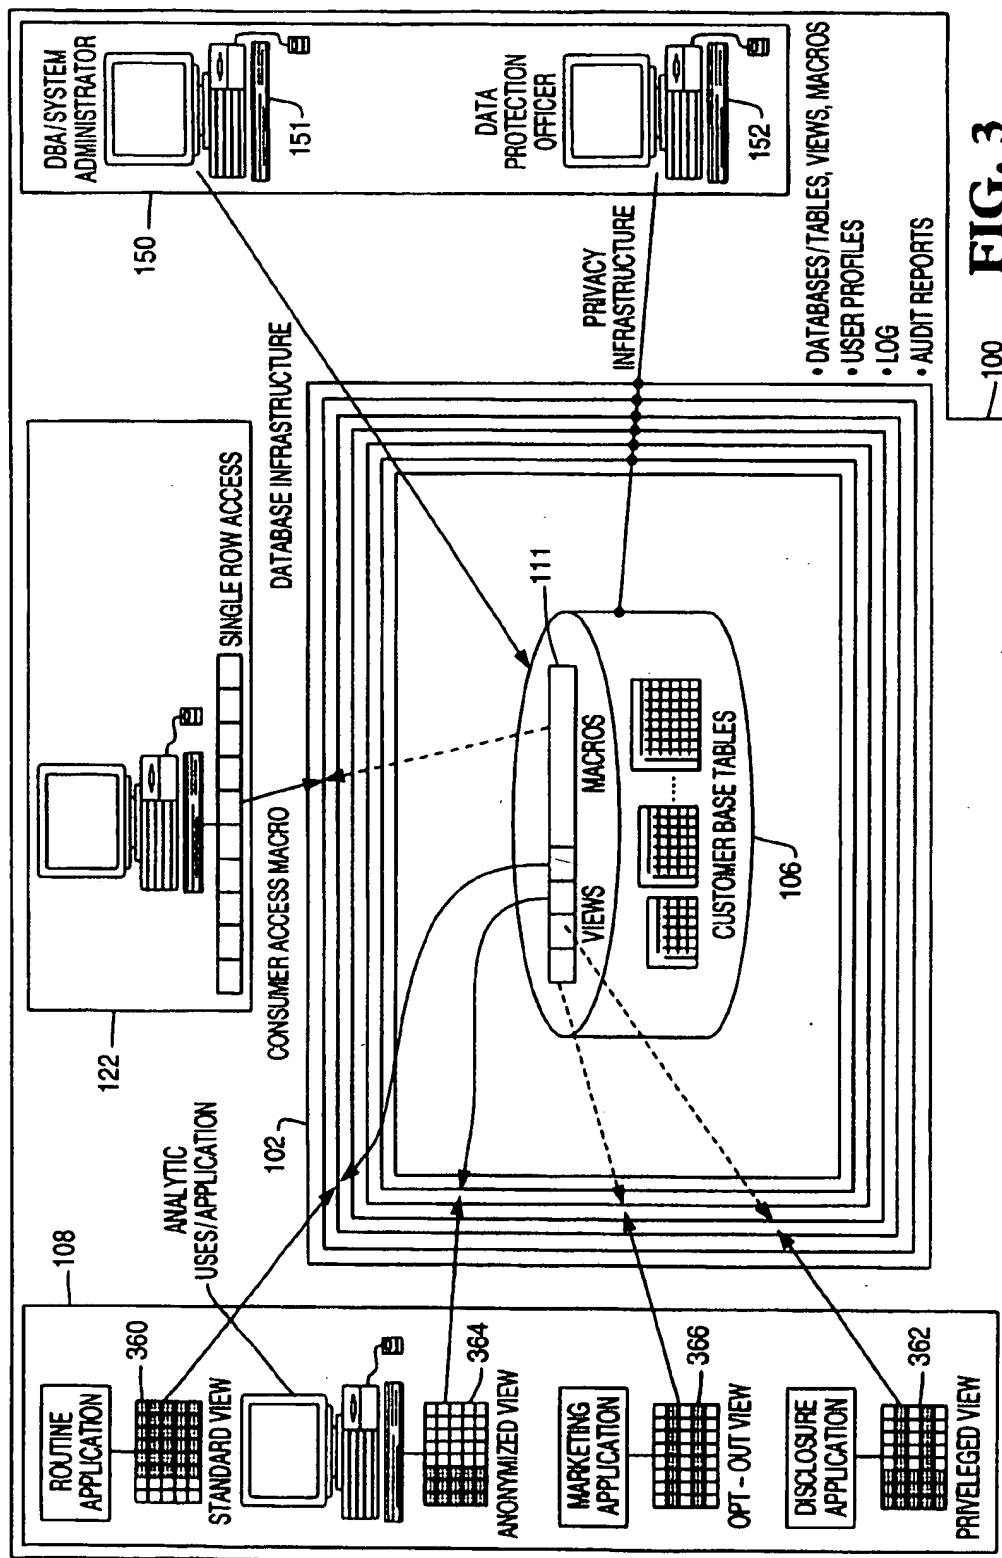


FIG. 3

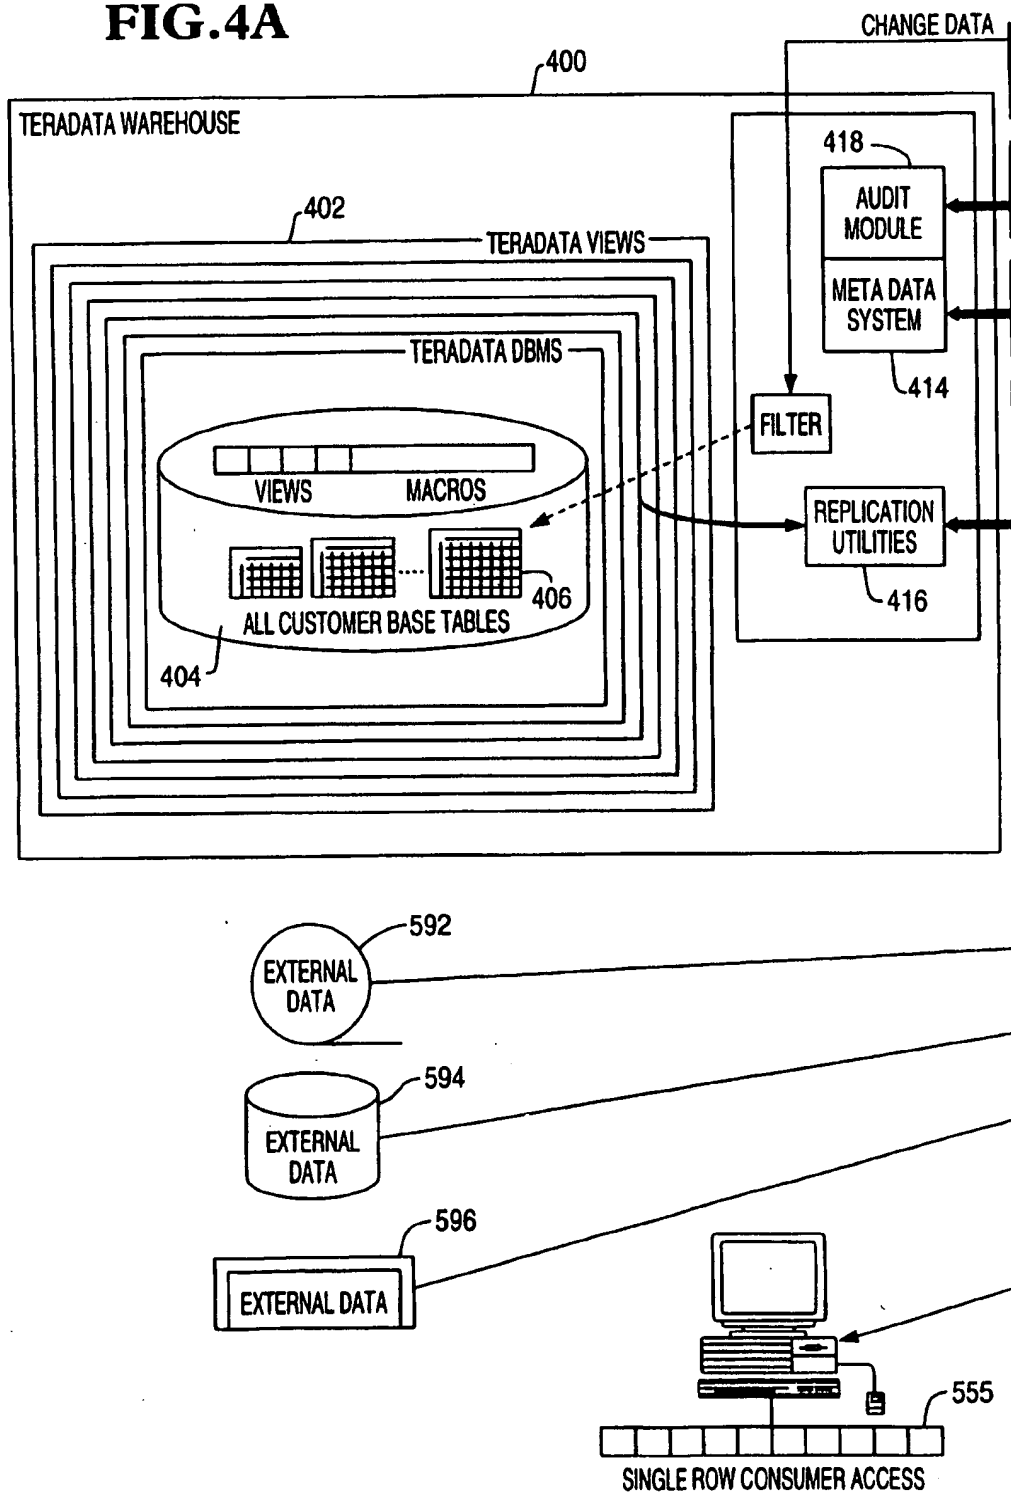
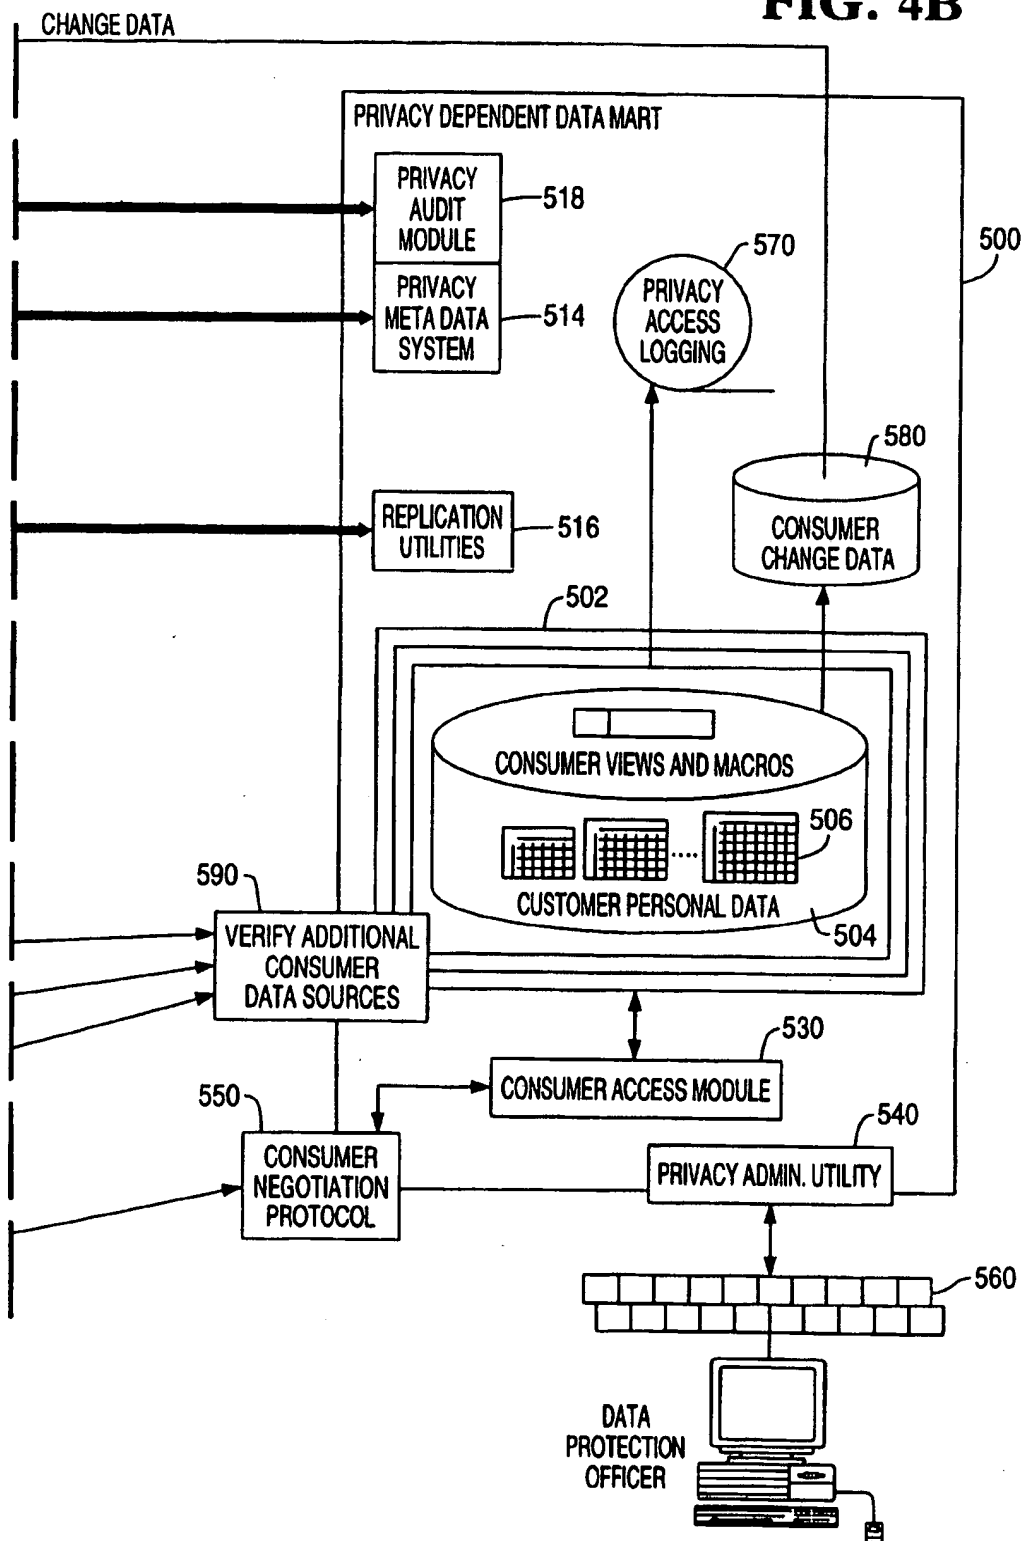
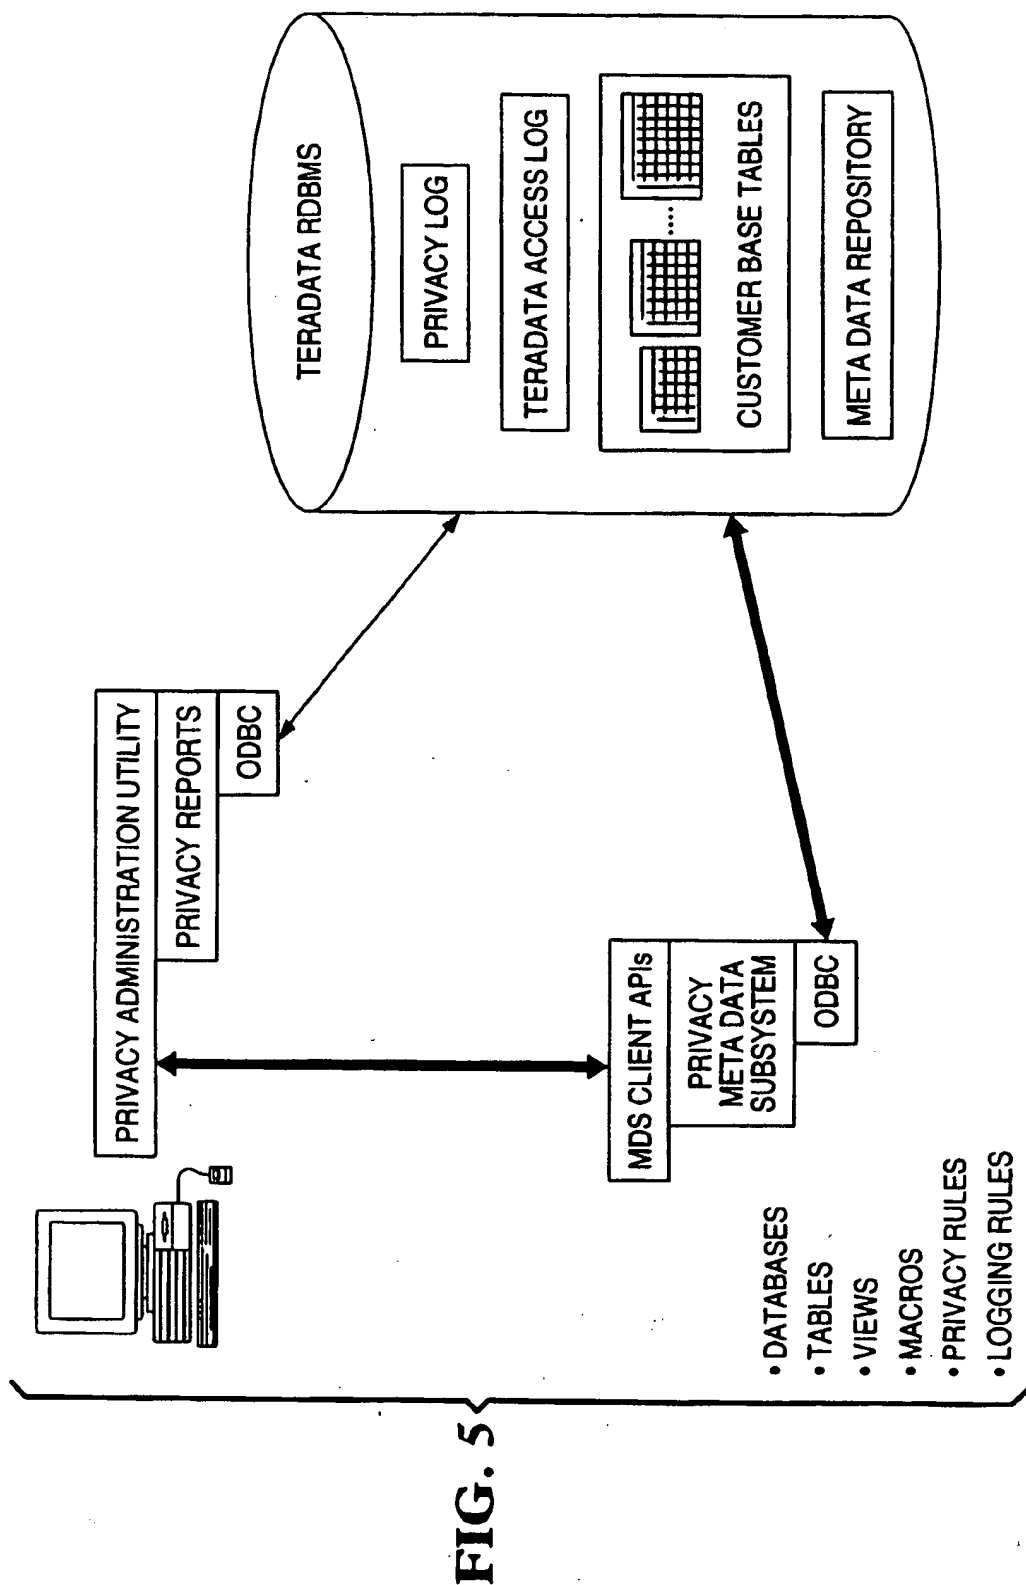
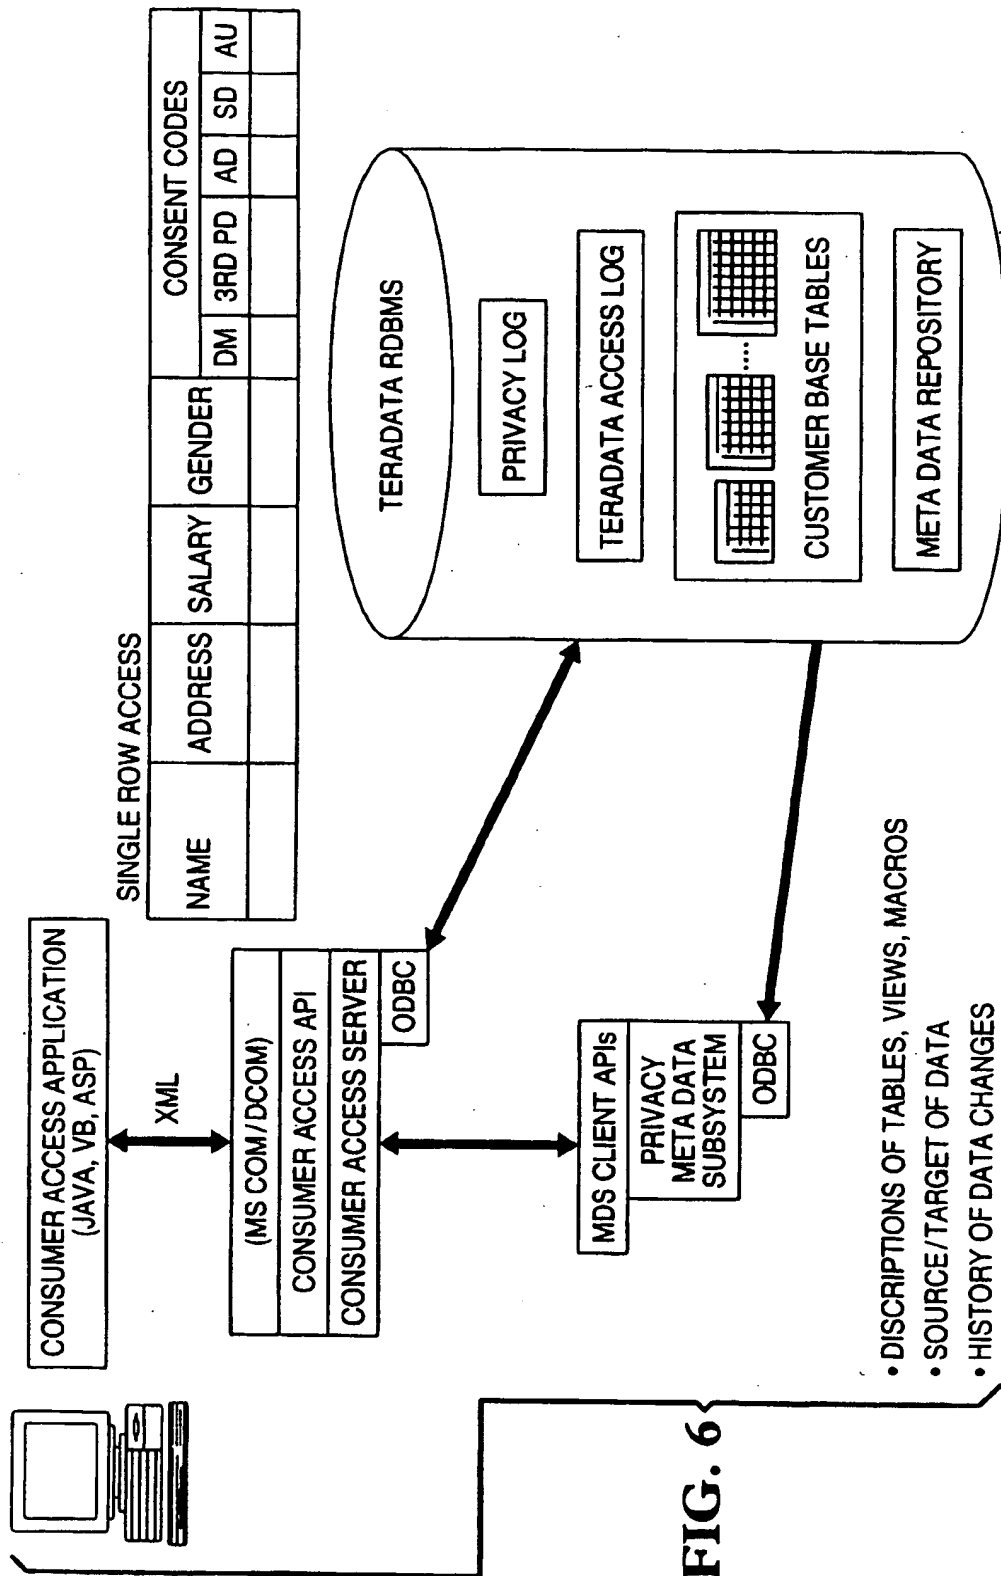
FIG. 4A

FIG. 4B







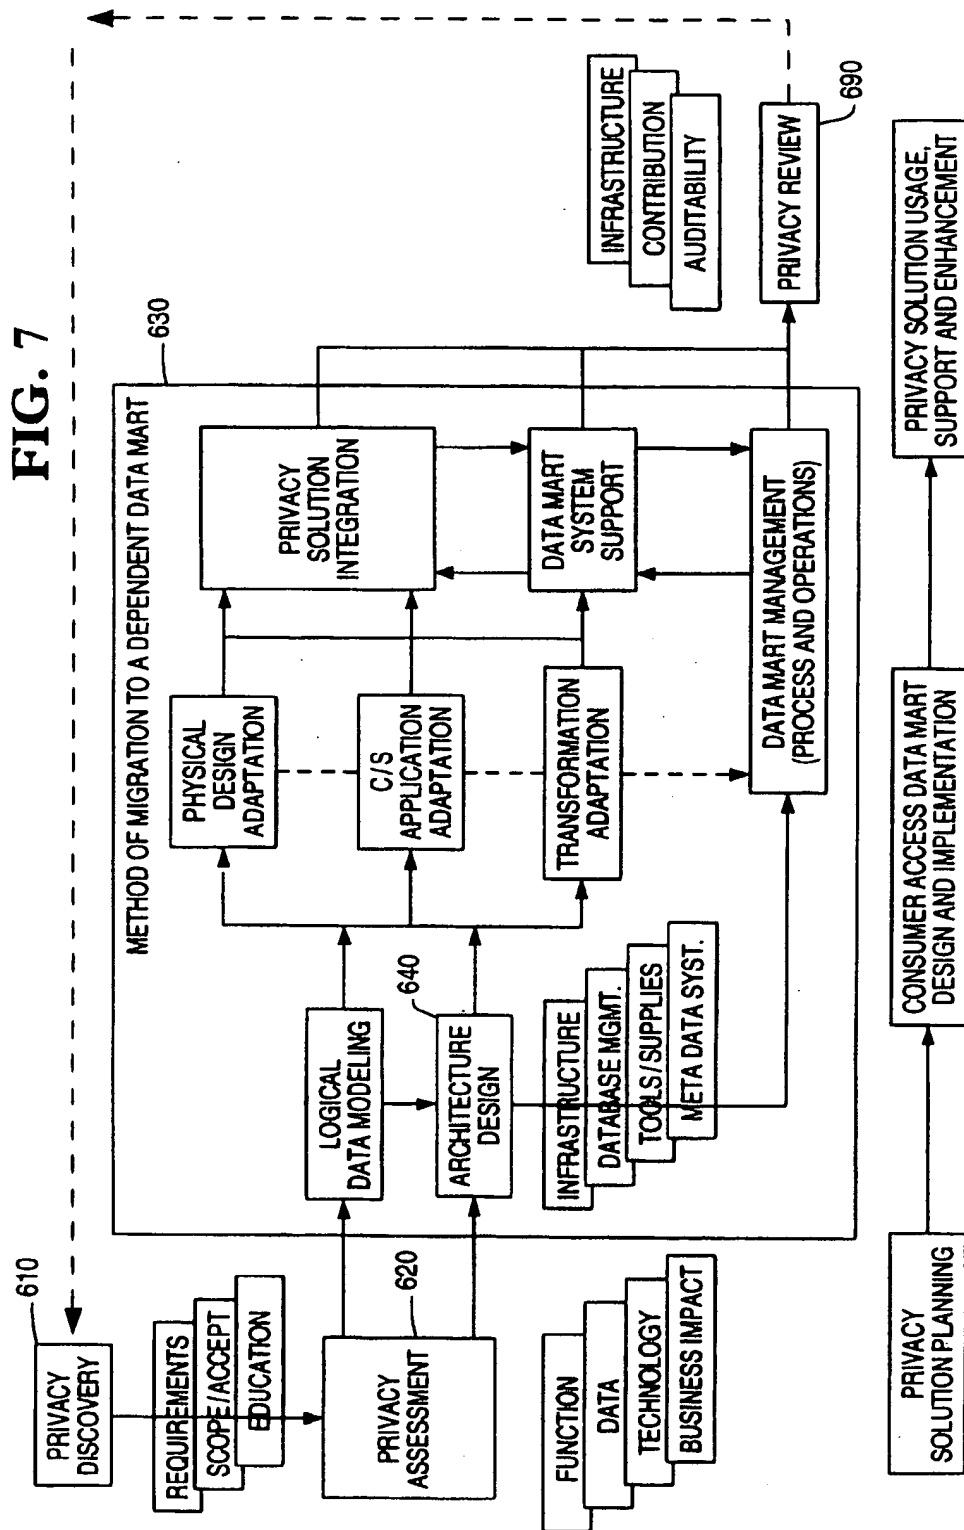


FIG. 8A

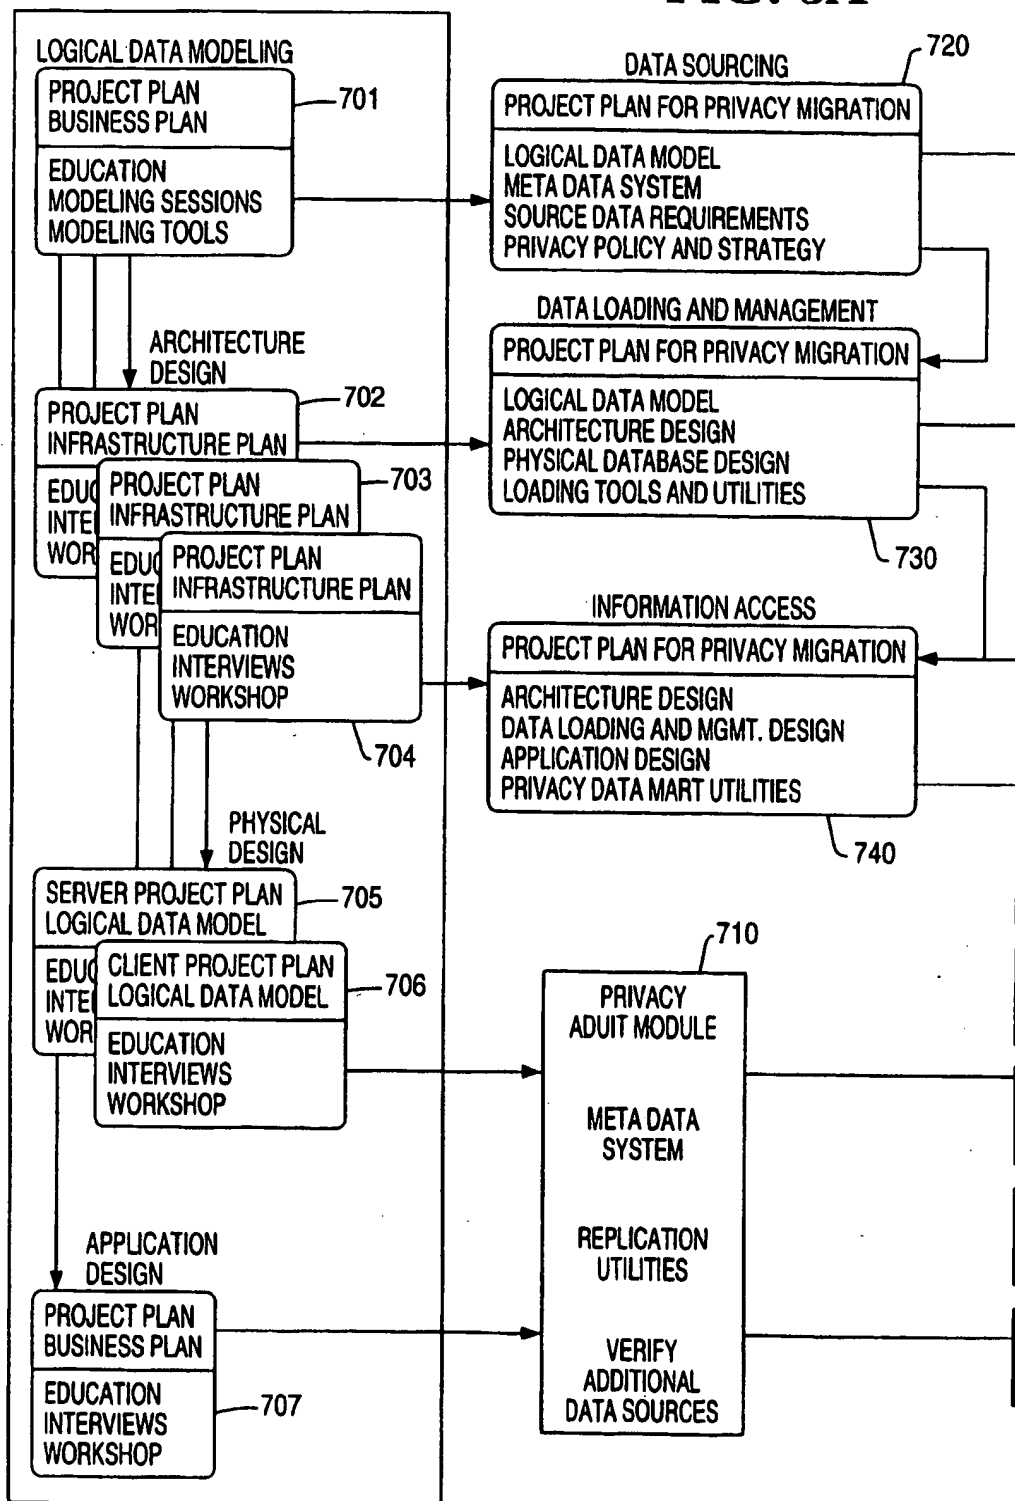
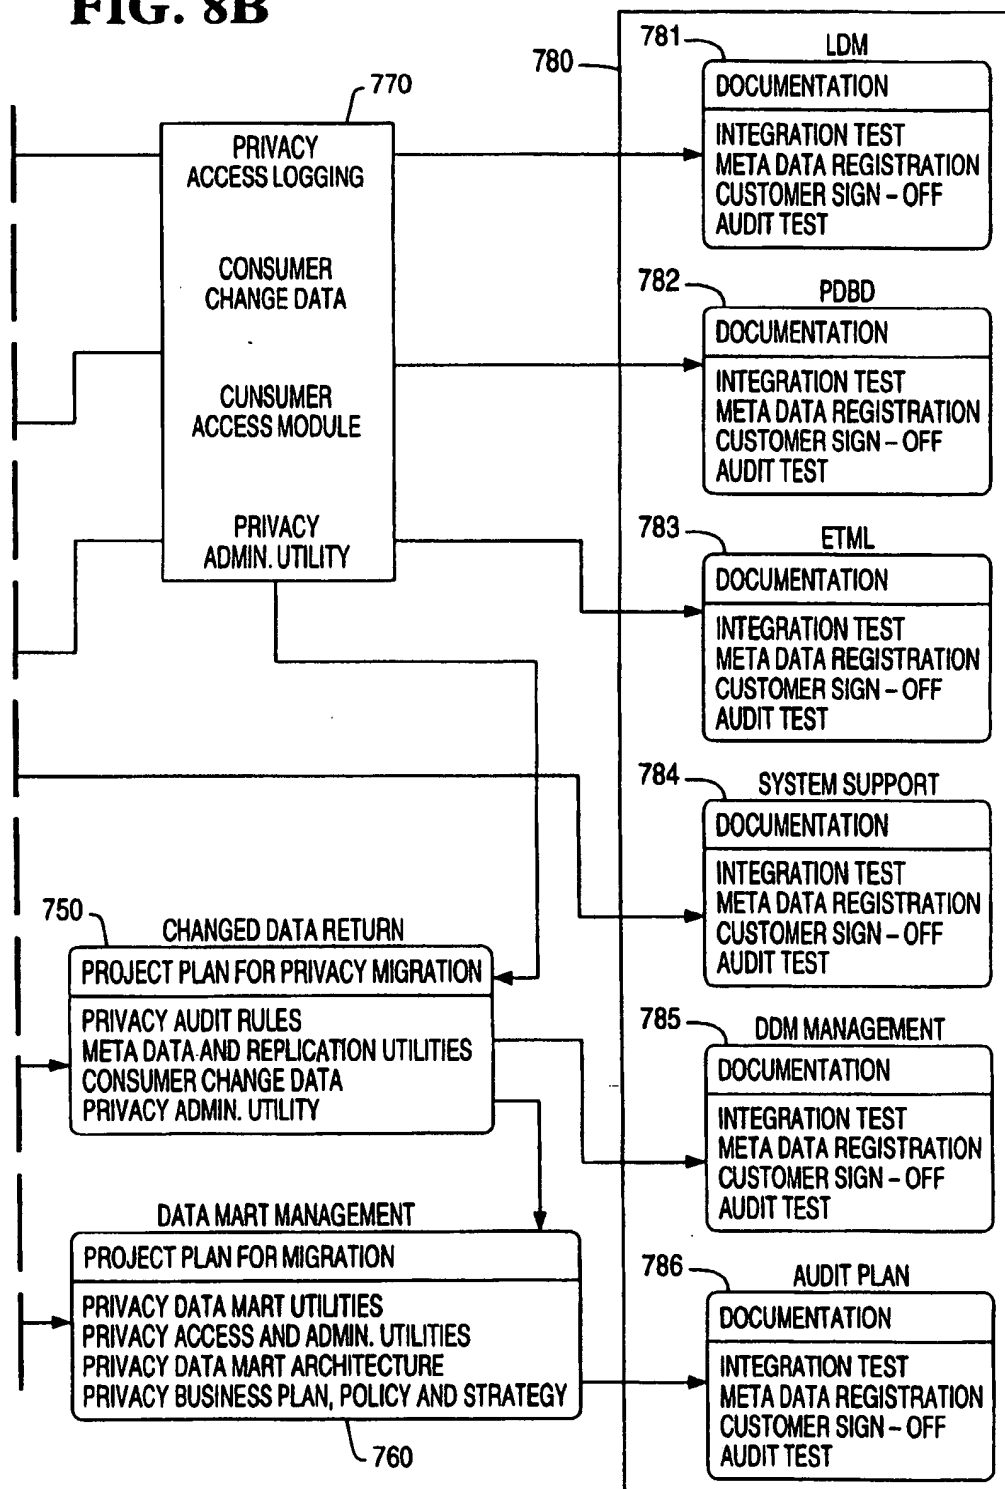


FIG. 8B



SYSTEM AND METHOD FOR MANAGING DATA PRIVACY IN A DATABASE MANAGEMENT SYSTEM INCLUDING A DEPENDENTLY CONNECTED PRIVACY DATA MART

This is a continuation-in-part of application Ser. No. 09/165,784, entitled "PRIVACY-ENHANCED DATABASE," by Kenneth W. O'Flaherty, Richard G. Stellwagen, Jr., Todd A. Walter, Reid M. Watts, David A. Ramsey, Adriaan W. Veldhuisen, Renda K. Ozden, and Patrick B. Dempster filed on Oct. 2, 1998, now U.S. Pat. No. 6,253,203.

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to the following co-pending and commonly assigned applications, each of which is hereby incorporated by reference herein:

U.S. patent application Ser. No. 09/165,457, entitled "PRIVACY-ENABLED LOYALTY CARD SYSTEM AND METHOD," by Kenneth W. O'Flaherty, Reid M. Watts, and David A. Ramsey, filed on Oct. 2, 1998; and

U.S. patent application Ser. No. 09/165,777 U.S. Pat. No. 6,275,824, entitled; "SYSTEM AND METHOD FOR MANAGING DATA PRIVACY IN A DATABASE MANAGEMENT SYSTEM," by Kenneth W. O'Flaherty, Richard G. Stellwagen, Jr., Todd A. Walter, Reid M. Watts, David A. Ramsey, Adriaan W. Veldhuisen, Renda K. Ozden, and Patrick B. Dempster filed on Oct. 2, 1998; and

U.S. Provisional Patent Application Serial No. 60/102, 832, entitled "SYSTEM AND METHOD FOR PRIVACY ENHANCED DATA WAREHOUSING," by Kenneth W. O'Flaherty, Richard G. Stellwagen, Jr., Todd A. Walter, Reid M. Watts, David A. Ramsey, Adriaan W. Veldhuisen, Renda K. Ozden, and Patrick B. Dempster filed on Oct. 2, 1998.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to systems and methods of data warehousing and analysis, and in particular to a system and method for providing consumer notification, access, data correction and change of preferences for data privacy in a data warehousing system that includes a physically separate but dependently connected data mart.

2. Description of the Related Art

Database management systems are used to collect, store, disseminate, and analyze data. These large-scale integrated database management systems provide an efficient, consistent, and secure data warehousing capability for storing, retrieving, and analyzing vast amounts of data. Meta Data Services are a comprehensive solution for managing metadata in complex data warehouse environments. Meta Data Services provides a solution for locating, consolidating, managing and navigating warehouse metadata. It also allows for setting aside an area from where all system aspects of privacy are registered, administered and logged in an auditable format. The ability to collect, analyze, and manage massive amounts of information through metadata has become a virtual necessity in business today, particularly when multiple hardware systems are involved.

The information stored by these data warehouses can come from a variety of sources. One important data ware-

housing application involves the collection and analysis of information collected in the course of commercial transactions between retailer outlets and retail consumers. For example, when an individual uses a credit card to purchase an item at a retail store, the identity of the customer, the item purchased, the purchase amount and other similar information are collected. Traditionally, this information is used by the retailer to determine if the transaction should be completed, and to control product inventory. Such data can also be used to determine temporal and geographical purchasing trends.

The data collected during such transactions is also useful in other applications. For example, information regarding a particular transaction can be correlated to personal information about the consumer (age, occupation, residential area, income, etc.) to generate statistical information. In some cases, this personal information can be broadly classified into two groups: information that reveals the identity of the consumer, and information that does not. Information that does not reveal the identity of the consumer is useful because it can be used to generate information about the purchasing proclivities of consumers with similar personal characteristics. Personal information that reveals the identity of the consumer can be used for a more focused and personalized marketing approach in which the purchasing habits of each individual consumer differentiates the approach and brings competitive advantage.

Unfortunately, while the collection and analysis of such data can be of great public benefit, it can also be the subject of considerable abuse. It can discourage the use of emerging technology, such as cash cards and loyalty card programs, and foster continuation of more conservative payment methods such as cash and checks. In fact, public concern over privacy is believed to be a factor holding back the anticipated explosive growth in web commerce.

For all of these reasons, when personal information is stored in data warehouses, it is incumbent on those that process and control this data to protect the data subjects from such abuse. As more and more data is collected in this, the computer age, the rights of individuals regarding the use of data pertaining to them have become of greater importance. What is needed is a system and method which provides all the advantages of a complete data warehousing system, while addressing the privacy concerns of the consumer. Consumers should have insight in what data about them is subject to collection and use.

Therefore, it is the responsibility of those that process and control personal data to provide accurate and full disclosure of what data is collected and processed, for what purposes, and under what limits of use. This includes data which the data controller has not collected directly from the consumer. It is the obligation of a data controller to provide access to the consumer of data which are being processed, in order to notify the consumer of the existence of a processing operation and, where data are collected from him, accurate and full information to verify in particular the accuracy of the data and the implied or explicitly stated preferences of privacy or data protection that has been agreed between the data controller and the data subject and work directly with the consumer to negotiate privacy preferences.

SUMMARY OF THE INVENTION

To address the requirements described above, the present invention discloses a method and apparatus for managing consumer notification and access and a means of correction and change of preferences for privacy or data protection in

a data warehousing system including a physically separate but dependently connected data mart.

The apparatus comprises a database management system, for storing data from a plurality of consumer database tables, with irrevocable logging of all access, whether granted or denied, to the data contents stored in the consumer data tables, a privacy metadata system that administers and records all data, users and usage of data that is registered as containing privacy elements, a replication system that feeds the consumer access system with personal consumer data, maintains integrity of the consumer data and provides changes and corrections back to the originating database management system through their own integrity filters as well as a means of storage and the mechanism to provide input for changes in the personal data or privacy preferences.

The method is supported by a privacy administrators utility and includes procedures for migration of consumer data from any state or format into a consistent and presentable state in the consumer access dependent data mart by establishing a database logical data model and physical database design in the data mart with all the tables, views and macros needed to reflect all aspects of personal data and its identifiers, dependently coupled for integrity to the base consumer database management system as a direct reflection of the tables in that system, extending database tables to store and retrieve privacy preference parameters for the data stored in the database table, the privacy parameters collectively reflected in a plurality of database views associated with the data, accepting personal data and privacy parameters from the data source, possibly including sources external to the data warehouse, storing the privacy parameters in the columns associated with the data, providing notification of and access to the data in the database table to a requesting consumer solely through a privacy metadata services interface in accordance with the personal privacy parameters.

Where possible the data models will be adapted to accepted privacy standards, like P3P, to reflect the data types and privacy sensitivity levels necessary and the consumer privacy preferences, provide for an adapted system for loading, formatting and maintaining data through Teradata utilities provide a system for returning changes back to the source system and a utility that allows a privacy administrator or data protection officer to manage the consumer access system to legal specifications. The program storage device comprises a medium for storing instructions performing the method steps outlined above.

BRIEF DESCRIPTION OF THE DRAWINGS

Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIG. 1 is a system block diagram of an exemplary embodiment of a data warehouse system;

FIGS. 2A and 2B illustrate a graphical representation of the privacy logical data model that supports the implementation of both the data warehouse and a dependent data mart;

FIG. 3 is a block diagram presenting an illustrative example of the structure of privacy-extended customer tables stored in the data management system and the database views that provide virtual separation between different user types and the actual data;

FIGS. 4A and 4B illustrate a data warehouse with a physically separate but dependently connected, privacy dependent data mart and the functions associated with the data mart;

FIG. 5 is a block diagram illustrating the functions of the privacy administration utility that supports the privacy dependent data mart.

FIG. 6 is a block diagrams illustrating the functions of the privacy consumer access module and utility that supports the privacy dependent data mart.

FIG. 7 is a flow chart illustrating the total methodology for building privacy into a data warehouse or a data mart consisting of a Privacy Planning phase, a Design & Implementation phase and a Privacy Usage, Support & Enhancement phase.

FIGS. 8A and 8B provide a graphical representation of the migration methodology that supports the implementation of the consumer access dependent data mart.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

In the following description, reference is made to the accompanying drawings which form a part hereof, and which is shown, by way of illustration, several embodiments of the present invention. It is understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

FIG. 1 is a system block diagram presenting an overview of a data warehousing system 100. The system comprises secure data warehouse 102 having a database management system 104 storing one or more extended databases 106 therein.

One important capability of a database management system is the ability to define virtual table and save that definition in the database as metadata with a user-defined name. The object formed by this operation is known as a dataview. As a virtual table, a dataview is not physically materialized anywhere in the database until it is needed. All accesses to data, other than for data administrative purposes, would be accomplished through dataviews. Various dataviews exist for purposes of implementing privacy rules. Metadata about the privacy dataview (including the dataview name, names and data types of the dataview columns, and the method by which the rows are to be derived) is stored persistently in the databases metadata, but the actual data presented by the view is not physically stored anywhere in association with the derived table. Instead, the data itself is stored in a persistent base table, and the view's rows are derived from that base table. Although the dataview is a virtual table, operations can be performed against dataviews just as they can be performed against the base tables.

The secure data warehouse 102 further comprises a suite of privacy metadata dataviews 108 through which all data in the extended database 106 are presented. Data within the extended database 106 can be viewed, processed, or altered only through the dataviews in this suite. The schema and logical model of the extended database and dataviews is set forth more fully herein with respect to FIG. 2.

Virtually all access to the data stored in the extended database 106 is provided solely through the dataview suite 108. Thus, retailer applications 110 and third party applications 112 have access only to such data as permitted by the database view provided. In one embodiment, provision is made to permit override of the customer's privacy preferences. However, in such circumstances, data describing the nature of the override is written to the database for retrieval by the audit module 118, so that the override cannot occur surreptitiously. Further, overrides may be monitored by the privacy metadata monitoring extensions 114 to provide an alert to the consumer when such overrides occur 116.

The limiting access to the data stored in the extended database 106 to access provided by the privacy dataview suite 108 for purposes of implementing privacy rules pro-

5

vides the capability to make the personal data anonymous (through the anonymizing view described herein), to restrict access to opted-out columns, which can apply to all personal data, separate categories of personal data, or individual data columns, and to exclude entire rows (customer records) for opted-out purposes—a row is excluded if any of the applicable opt-out flags is on for the customer in question.

Using a client interface module 122 that communicates with the dataviews 108, a client 124 can access, control, and manage the data collected from the client 124. This data control and management can be accomplished using a wide variety of communication media 140, including the Internet 126 (via a suitable browser plug-in 125, a modem 130, voice telephone communications 132, or a kiosk 134 or other device at the point of sale. To facilitate such communications, the kiosk or other device at the point of sale, can issue a smartcard 136 or a loyalty card 138. The kiosk/pos device 134 can accept consumer input regarding privacy preferences, and issue a smartcard 136 or loyalty card 138 storing information regarding these preferences. Similarly, when using the kiosk/pos device 134 and the smartcard 136 or loyalty card 138, the consumer may update or change preferences as desired. In cases where the loyalty card 138 is a simple read only device (such as a bar-coded attachment to a key ring), the kiosk/pos device 134 can accept issue replacement cards with the updated information as necessary. Transactions using the loyalty card 138 or smartcard 136 are selectively encrypted. Either card may interact directly with the server or through a plug-in to implement the security rules selected.

Through this interface, the consumer can specify data sharing and retention preferences. These allow the consumer to specify when and under what circumstances personal information may be retained or shared with others. For example, the consumer may permit such data retention as a part of a loyalty card program or specify that use of the data is limited to particular uses. Further, the consumer may specify under what circumstances the data may be sold outright, used for statistical analysis purposes, or used for selective marketing programs.

The data warehousing system 100 also permits use of anonymous data within the data warehouse 102 via a privacy service 150. When the user desires anonymous data, the transaction is routed to the privacy service 150. The privacy service 150 accesses a privacy rule database 152 and other security information 154 and uses the privacy rule and security information to remove all information from which the identity of the consumer can be determined. The cleansed transaction information response is then forwarded to the anonymity protection interface module 160 in the secure data warehouse. Communications with the secure data warehouse 102 use a proxy user identification, which is created by the privacy service 150 from the customer's username or other identifying information. If the customer does not require anonymous data, the transaction is provided directly to the retailer who may store the transaction information response in the extended database.

Since it alone provides access to data within the extended database, the dataview suite 108 also provides a convenient and comprehensive means for auditing the security of the secure data warehouse 102.

The secure data warehouse 102 also comprises metadata monitoring extension 114. This extension 114 allows the customer to generate a rule to monitor the use of personal data, and to transmit an alert 116 or callback if a metadata definition change occurs. The customer can control the

6

metadata monitoring extension 114 to trigger an alert when the consumer's personal information is read from the extended database 106, when personal information is written to the extended database 106, when opt-out delimiters stored in the extended database are changed, or when a table or a dataview is accessed. The metadata monitoring extension 114 also records data source information, so customers can determine the source of the data stored in the secure data warehouse 102. The data source may be the customer, or may be a third party intermediary source. This feature is particularly useful when the consumer would like to not only correct erroneous information, but to determine the source of the erroneous information so the error will not be replicated in the same database or elsewhere.

The metadata monitoring extension 114 can also be used to support auditing functions by tracking reads or writes from the extended database 106 as well as the changes to the dataview suite 108.

The described system can be implemented in a computer comprising a processor and a memory, such as a random access memory (RAM). Such computer is typically operatively coupled to a display, which presents images such as windows to the user on a graphical user interface. The computer may be coupled to other devices, such as a keyboard, a mouse device, a printer, etc. Of course, those skilled in the art will recognize that any combination of the above components, or any number of different components, peripherals, and other devices, may be used with the computer.

Generally, the computer operates under control of an operating system stored in the memory, and interfaces with the user to accept inputs and commands and to present results through a graphical user interface (GUI) module. Although the GUI module is typically a separate module, the instructions performing the GUI functions can be resident or distributed in the operating system, an application program, or implemented with special purpose memory and processors. The computer may also implement a compiler that allows an application program written in a programming language such as COBOL, C++, FORTRAN, or other language to be translated into processor-readable code. After completion, the application accesses and manipulates data stored in the memory of the computer using the relationships and logic that was generated using the compiler.

In one embodiment, instructions implementing the operating system, the computer program, and the compiler are tangibly embodied in a computer-readable medium, e.g., data storage device 170, which could include one or more fixed or removable data storage devices, such as a zip drive, floppy disc drive, hard drive, CD-ROM drive, tape drive, etc. Further, the operating system and the computer program are comprised of instructions which, when read and executed by the computer, causes the computer to perform the steps necessary to implement and/or use the present invention. Computer program and/or operating instructions may also be tangibly embodied in memory and/or data communications devices, thereby making a computer program product or article of manufacture according to the invention. As such, the terms "program storage device," "article of manufacture" and "computer program product" as used herein are intended to encompass a computer program accessible from any computer readable device or media.

Those skilled in the art will recognize many modifications may be made to this configuration without departing from the scope of the present invention. For example, those skilled in the art will recognize that any combination of the

above components, or any number of different components, peripherals, and other devices, may be used with the present invention.

FIGS. 2A and 2B provide a diagram showing the logical model of the secure data warehouse 102 and the dataview suite 108 in greater detail. The extended database 106 comprises a customer table 200, which is segmented into categories of personal data: such as phone 218, address 216, demographic 202, employer 204, financial account 210, navigation history 214, transaction history 206, and online contact 208. Each personal data category also has an associated consent table: such as phone consent 238, address consent 234, demographic consent 230, employer consent 220, financial consent 228, navigation consent 232, transaction consent 226, and online contact consent 224. The consent tables specify data reflecting the privacy preferences, or "opt-outs", for the accompanying data. In the disclosed embodiment, these privacy preferences include "opt-outs" for (1) direct marketing 240, (2) disclosure of personal data along with information identifying the consumer 242, (3) anonymous disclosure of personal data 242, (4) disclosure of personal data for purposes of making automated decisions 244, and (5) disclosure or use of sensitive data 246. Start and end dates are also maintained within the consent tables for historical tracking of consumer consent options.

In the logical data model, the individual consent tables allow very fine-grained selection by the consumer of privacy preferences. For example, the consumer could opt-in to third party disclosure of her phone number, but opt-out to third party disclosure of her address. The model also allows privacy preferences that apply across the entire consumer record, store in the privacy consent codes table 236. The automated decision code 244 allows consumers to indicate whether their data could be used to perform automated processing. The sensitive data code 246 allows consumers to permit dissemination of sensitive data.

In one embodiment, an NCR Corporation TERADATA database management system is utilized to implement the foregoing logical model. This implementation has several advantages.

First, the TERADATA database management system's ability to store and handle large amounts of data eases the construction of the many different views and allows the secure data warehousing system 100 to utilize a logical data model is in or close to the third normal form.

Second, unlike systems which execute SQL queries as a series of selections to narrow the data down to the dataview subset, the TERADATA database management system rewrites dataview-based queries to generate the SQL that selects the necessary columns directly from the appropriate base tables. While other views materialize entire tables before narrowing down the data to the view subset, TERADATA generates SQL that selectively pulls appropriate columns and rows into the result table. This method is a particularly advantageous in implementing the foregoing logical model.

Third, the foregoing logical model generally results in dataviews, which include complex queries and wide SQL expressions. The TERADATA database management system is particularly effective at optimizing such queries and SQL expressions.

FIG. 3 illustrates a number of dataviews that are provided in the dataview suite 108. These dataviews include a standard view 360, a privileged view 362, an anonymizing view 364, and an opt-out view 366. These views limit visibility

into the data in the customer table 106 in accordance with the values placed in the data control columns.

The standard view 360 will not present personal data unless either the flag in column (indicating that the personal information and identifying information can be disseminated or indicating that personal information can only be disseminated anonymously) is activated. Hence, the standard view 360 selectively masks personal data from view unless the consumer has set the appropriate flags to the proper value.

Scaleable data warehouse (SDW) customer Data Base Administrator's (DBA) 151 set up views into customer tables (any tables containing personal information about their customers), controlled by the Data Protection Offices 152, such that, for routine users, all columns of personal information are hidden.

~~This interface module 122, which is used to view, specify, and change consumer privacy preferences, is a privileged application. Appropriate security measures are undertaken to ensure that the privileged applications are suitably identified as such, and to prevent privileged view access by any entity that is not so authorized.~~

Certain SDW applications ("Class B") may perform analysis on personal data, in order to gain insight into customer behavior, e.g. to identify trends or patterns. Such applications may be driven by end-users (knowledge workers or "power analysts") performing "ad hoc" queries, typically using either custom-built software or standard query or OLAP Tools, where the end-user spots the patterns. They may also involve the use of data mining tools, where statistical or machine learning algorithms, in conjunction with the analyst, discover patterns and from them build predictive models.

FIGS. 4A and 4B illustrate a data warehouse 400 with a physically separate but dependently connected, privacy dependent data mart 500 and the functions associated with the data mart. The data warehouse includes a data base management system 404 storing one or more database tables 406 containing personal data 408. Communication between the data warehouse 400 and the privacy dependent data mart is provided through audit 418, metadata system 414, and replication 416 modules contained within data warehouse 400 and corresponding privacy audit 518, privacy metadata system 514, and replication 516 modules contained within privacy dependent data mart 500. In this embodiment, each class of functionality is applied separately to the data (e.g. filtering the change data), including specific control functions (e.g. providing audit reports or replicating data). For example, the data warehouse 400 contains the only version of all consumer information; all changes to the structure and use are fully audited and all input to the data contents or consumer preferences are filtered and limited for integrity. These limitations can be selected by entering the proper combination of integrity and preference. The present invention permits the expansion of the above described privacy preference paradigm to a similar system of multiple functions of consumer information and preferences, based upon the same detail of customer preferences.

In the privacy dependent data mart embodiment, the security and privacy protection features of the extended database are further enhanced with the use of privacy access logging 570 that captures all access attempts to the customer data, whether granted or denied, and the consumer change data 580 as provided by the customer that examines their own data and preferences. This may be used by the system on-line or in batch mode to feed the authorized changes back to the source database through integrity filters.

In one embodiment, external data in various formats 592, 594 and 596 might be allowed to enrich the consumer data 590 through an additional privacy data source filter, and selectively applied to the consumer personal data. This technique allows external customers data to be automatically flagged (e.g. for authentication purposes), but could allow for exclusion of processing for return of change data back to the data warehouse.

FIG. 5 is a block diagram illustrating the functions of the privacy administration utility 540 that supports the privacy dependent data mart.

FIG. 6 is a block diagram illustrating the functions of the privacy consumer access module 530 that supports the privacy dependent data mart.

FIG. 7 is a flow chart illustrating the total methodology for building privacy into a data warehouse or a data mart consisting of a Solution Planning phase, a Design and Implementation phase and a Solution Usage, Support and Enhancement phase. The functions of the Privacy Discovery service 610 are to provide education, determine the business requirements, and set the scope to be accepted by the business. Privacy Assessment service 620 is based on the outcome of Privacy Discovery and executes a GAP analysis against the functional, data, and technical requirements for Privacy and uses these evaluations as input for the Business Impact Assessment which quantifies the impact that implementation choices will bring to the current business in terms of investment and revenue opportunity, positive or negative. Privacy Assessment also creates an implementation blueprint of the changes needed in infrastructure and business practices to enable a data warehouse for Privacy. This blueprint feeds into the Architecture Design 640 that lays the foundation for choices for change in Infrastructure, Database Management, Tools and Utilities all built around an integrated Metadata system. After completion of an implementation of Privacy in a data warehouse environment a Privacy Review 690 is recommended to evaluate whether the implementation goals for infrastructure change has been met and what Data Warehouse Contributions have been achieved. This service also prepares for auditability by EDP Auditors or Privacy or Data Protection regulators.

FIG. 8 is a flow chart illustrating the specific methodology for building the Consumer Access Dependent Data Mart and migrating consumer data and its accompanying profile for privacy preferences from a data warehouse and other data sources to the data mart.

Project Management—Project Management is critical to the success of Dependent Data Mart Migration to meet obligations to the customer and for the elimination of 'scope creep', a project plan is required for all implementations. The Project Plan governs the Design Phase 700 with Logical Data Modeling 701, Architecture Design 702 (Source data), 703 (Target Data) and 704 (Data Mart), Physical Design 705 (Business Profile) and 706 (Consumer Profile) and Application Design 707. Each step in the Design Phase contains Education, Interview and Workshop elements that accompany the tasks necessary to complete the input into the next phase. Also, Logical Data Modeling 701 feeds information into Architecture Design 702, Physical Design 703 and Application Design 704.

Project Management also passes the plan from the Design steps to the Implementation services for Data Sourcing 720, Data Loading and Management Access 730, Information Access 740, Changed Data Return 750 and Data Mart Management 760. The NCR project management methodology is the single point of contact with the customer. Project managers are responsible for all aspects of the Dependent Data Mart program.

Logical Data Modeling—This service produces the attributed logical data model and/or star schema for the initial implementation of the Dependent Data Mart. Activities in this service include confirmation of requirements and generation of the data model showing relationships and attributes. The data model is crucial to a Dependent Data Mart solution to ensure that the proper business focus and flexibility are maintained in the solution. The data model is not specific to a platform or database and is separate from any physical dependencies. The data model for the Dependent Data Mart may be either a logical data model derived from the enterprise data warehouse, or a star schema data model.

Architecture Design—This service produces the infrastructure for the initial implementation of the Dependent Data Mart. Activities in this service include confirmation of requirements and generation of the source systems that feed the Dependent Data Mart, the Dependent Data Mart itself and the architecture for the return of changed data back to the data warehouse. The architecture model is crucial to a Dependent Data Mart solution to ensure that the proper technical focus and flexibility are maintained in the solution. The architecture model is specific to a platform and database and is based on its physical dependencies.

Physical Database Design—This service provides the client a physical database design optimized for dependent data mart. The primary activities of this service are: translating the data model to a physical database design, database construction, design optimization, and functional testing of the constructed database.

Application Design (Query Development)—This service provides the design and implementation of the query interface for the Dependent Data Mart Solution. Utilizing a GUI based tool, queries to answers of agreed upon business questions will be developed as part of the Dependent Data Mart Program. The Application Design service develops applications that enable review and input for change based on access to detail consumer data, data summaries, and staged queries.

Data Transformation and Replication—This service designs the process and develops the utilities and programming that allow the dependent data mart database to be initially loaded and maintained. The service locates, transforms, replications, transports, and loads data onto the target platform. Included is the operational planning that allows the reloading or incremental loading of the dependent data mart on a periodic basis. Data transformation and replication for the Dependent Data Mart Program will normally be executed using Teradata utilities.

Data Mart Management—This service encompasses the backup, archive, restore, and recovery strategy for the dependent data mart. This service does not include taking the dependent data mart into production, this is the responsibility of the Customer.

Documentation—This service encompasses the Integration Test, Meta Data Registration, Audit Testing and Customer sign-off. Customer Education is key to any data warehouse or dependent data mart success and is included as part of the dependent data mart services program. Other, standard Data Warehouse Implementation services elements are:

- Logical Data Model
- Physical Data Base Design
- Extract, Transfer, Move and Load scripts
- System Management Integration
- Audit and Control Plan

11

There are many types and uses of metadata including: Business rules and definitions, Directory of warehouse users, developers, users, etc., Database schema's and views, Transformational mappings, Source database logical models, Target warehouse models including data marts, Refresh frequency of data, Security, Reports, Performance metrics, and Computing system components. Thus, the content of metadata is evolved during Privacy Implementation from merely a logical model of the source and target databases to full integration with business rules to information about information system resources.

The foregoing description of the various embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many alternatives, modifications, and variations will be apparent to those skilled in the art of the above teaching. Accordingly, this invention is intended to embrace all alternatives, modifications, and variations that have been discussed herein, and others that fall within the spirit and broad scope of the claims.

What is claimed is:

1. A data warehousing, management, and privacy control system, comprising:

- a database management system, for storing and retrieving customer data;
- a privacy metadata system that administers and records all customer personal data, users of said customer personal data, and usage of said customer personal data;
- a replication system providing communication between said database management system and said privacy metadata system; and

12

a database management system interface operatively coupled to the database management system and controlling access to said customer data and to said customer personal data through said replication system.

2. The data warehousing, management, and privacy control system according to claim 1, wherein:

said replication system provides customer personal data from said database management system interface to said privacy metadata system.

3. The data warehousing, management, and privacy control system according to claim 1, further comprising:

a customer access module operatively coupled to the privacy metadata system and providing a customer with means to access data, correct data and change of preferences to customer personal data related to said customer.

4. The data warehousing, management, and privacy control system according to claim 1, wherein:

said replication system provides changes and corrections to said customer data from said privacy metadata system to said database management system.

5. The data warehousing, management, and privacy control system according to claim 1, wherein:

said database management system interface provides access to said customer data and to said customer personal data in accordance with privacy parameters stored in said database management system.

6. The data warehousing, management, and privacy control system according to claim 1, further comprising:

a privacy access logging system that captures and records all access attempts to said customer personal data.

* * * * *

United States Patent [19]

Balgeman et al.

US005446880A

[11] Patent Number: 5,446,880

[45] Date of Patent: Aug. 29, 1995

[54] DATABASE COMMUNICATION SYSTEM THAT PROVIDES AUTOMATIC FORMAT TRANSLATION AND TRANSMISSION OF RECORDS WHEN THE OWNER IDENTIFIED FOR THE RECORD IS CHANGED

[75] Inventors: Timothy E. Balgeman, Warrenville; Diane M. Clodi, Oswego; Robert W. Haddleton, Jr., Naperville; John R. North, Geneva; Judith A. Selby, Wheaton, all of Ill.

[73] Assignee: AT&T Corp., Murray Hill, N.J.

[21] Appl. No.: 937,814

[22] Filed: Aug. 31, 1992

[51] Int. Cl.⁶ G06F 17/30

[52] U.S. Cl. 395/600; 395/200.01; 364/241.7; 364/242.5; 364/284.3; 364/DIG. 1

[58] Field of Search 395/600, 200; 340/825.05; 364/DIG. 1, DIG. 2, 241.7, 241.8, 242.5, 384.3

[56] References Cited

U.S. PATENT DOCUMENTS

4,031,512	6/1977	Fabor	340/825.05
4,135,240	1/1979	Ritchie	395/600
4,714,995	12/1987	Materna et al.	395/600
4,714,995	12/1987	Materna et al.	395/600
4,933,846	6/1990	Humphrey et al.	395/325
5,119,465	6/1992	Jack et al.	395/500
5,249,231	9/1993	Covey et al.	395/425
5,276,869	1/1994	Forrest et al.	395/600
5,283,887	2/1994	Zachery	395/500
5,329,618	7/1994	Moati et al.	395/200

FOREIGN PATENT DOCUMENTS

0216535A2 1/1987 European Pat. Off. .

OTHER PUBLICATIONS

Papasoglou et al. "Distributed Heterogeneous Informa-

tion Systems & Schema Transformation", *Parbase-90, Conf Date, 7-9 Mar. 1990, pp. 388-397 IEEE*.

Pu, "Superdatabases for Composition of Heterogeneous Databases", *Pr. 4th Int. Conf. On Data Engineering*, date 1-5 Feb. 1988; pp. 548-555.

Chung, "DATAPLEX: an access to Heterogeneous distributed databases", *Communications Of The ACM*, Jan. 1990: V33 issue h1, p. 70-81.

J. Gray & S. Metz "Solving The Problems of Distributed Databases", *Data Communications*, vol. 12, No. 10, Oct., 1983, pp. 183-192.

R. Ahmed et al. "The Pegasus Heterogeneous Multidatabase Systems", *Computer*, vol. 24, No. 12, Dec., 1991, pp. 19-26.

"Financial Systems Interface", *IBM Technical Disclosure Bulletin*, vol. 31, No. 5, Oct., 1988, pp. 36-39.

Primary Examiner—Thomas G. Black

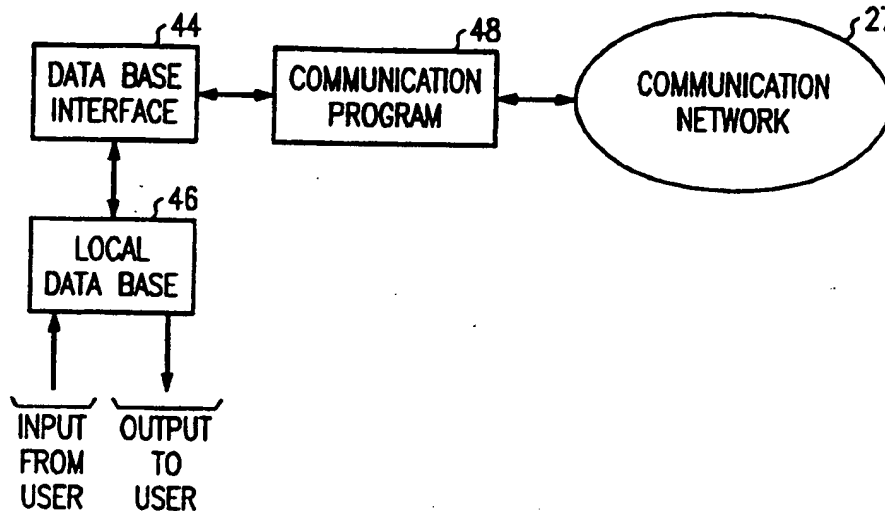
Assistant Examiner—Jack M. Choules

Attorney, Agent, or Firm—Charles L. Warren

[57] ABSTRACT

A database interface and associated communication program are installed in each node in a communication network. The database interface provides a translation from the record format of the database and a standardized format for transmission to other nodes thus providing translation between the formats of the different databases formats as records are transmitted between the databases on different nodes. The communication program transmits records that have been translated to the standardized format to nodes which contain corresponding records maintaining current records at each node. A node having a corresponding node can assign an owner. Assigning an owner causes the record to be sent to the database that is associated with the new owner. The owner has the responsibility for acting on the record. The originator of the record is also identified.

15 Claims, 6 Drawing Sheets



NODE	12	22	26	12	22	26	12	22	26
TICKET	UN12-1487			UN12-1487	UN12-1487		UN12-1487	UN12-1487	UN12-1487
ORIGINATOR	UN12			UN12	UN12		UN12	UN12	UN12
OWNER	UN12			N22	N22		N26	N26	N26
STATUS	OPEN			OPEN	OPEN		OPEN	OPEN	OPEN
T0			T1			T2			

FIG. 6

NODE	12	22	26	12	22	26
TICKET	UN12-1487	UN12-1487	UN12-1487	UN12-1487	UN12-1487	UN12-1487
ORIGINATOR	UN12	UN12	UN12	UN12	UN12	UN12
OWNER	UN12	UN12	UN12	UN12	UN12	UN12
STATUS	SOLVED	SOLVED	SOLVED	CLOSED	CLOSED	CLOSED
T3			T4			

FIG. 7

FIG. 8

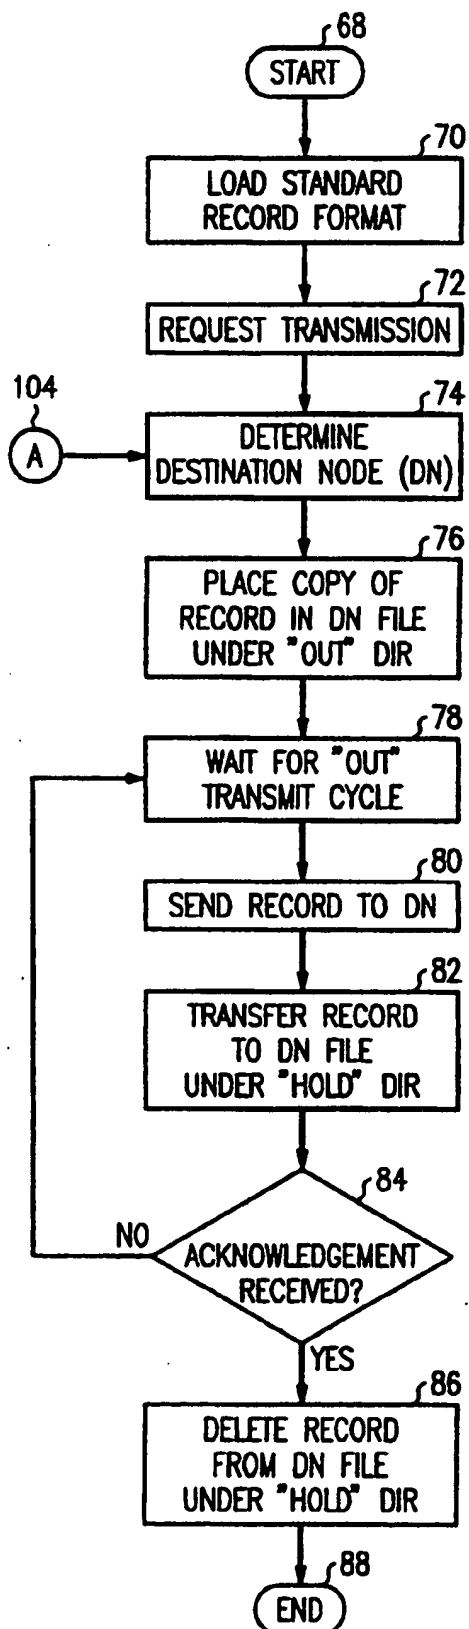
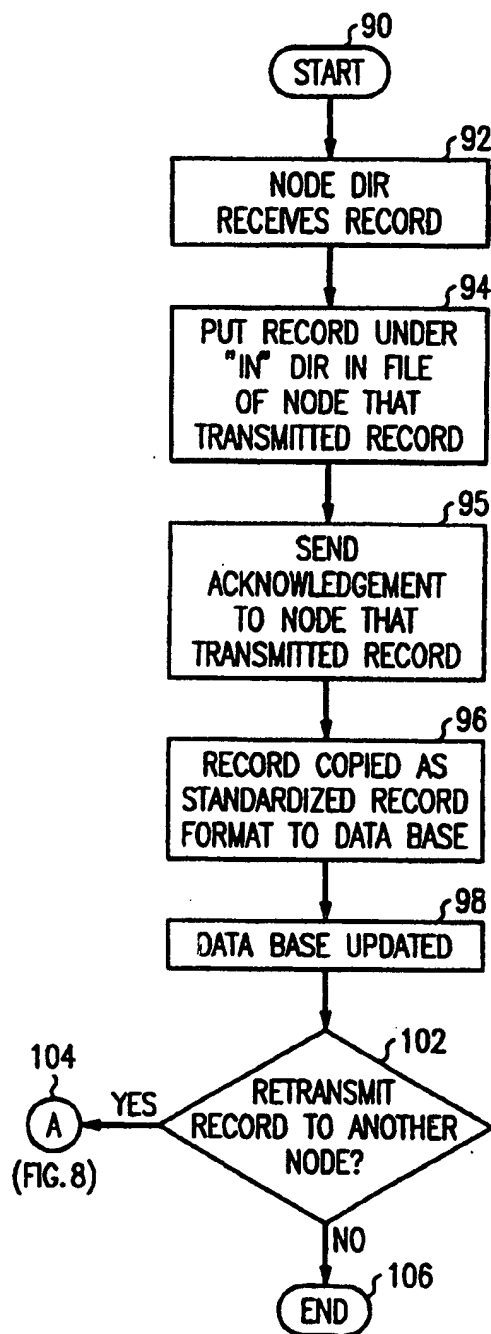


FIG. 9



DATABASE COMMUNICATION SYSTEM THAT PROVIDES AUTOMATIC FORMAT TRANSLATION AND TRANSMISSION OF RECORDS WHEN THE OWNER IDENTIFIED FOR THE RECORD IS CHANGED

BACKGROUND OF THE INVENTION

This invention relates to a communication system in which records are stored in independent databases and more specifically relates to communicating and updating corresponding records in the different independent databases. This invention is especially, but not exclusively, suited for intercompany database record sharing such as between an equipment manufacturer and its customer.

A variety of communication networks exist which can transmit data files from an origination node (NODE O) to a destination node (NODE D). However, such a transmission may not achieve the goal of transferring the desired information from NODE O to NODE D in a form readily usable at NODE D. As an example, assume that the information to be transmitted by NODE O is a record created in NODE O's independent, local database. If the objective is to create a duplicate record in an independent, local database at NODE D, the received record can be directly entered in NODE D's database assuming both NODE O and NODE D utilize the same database and the same field formats. If the database at NODE D is not directly compatible with the record received from NODE O, a user at NODE D must manually translate the received record into a record in the local database format. It is apparent that substantial inefficiencies exist where different database formats are utilized. If the databases utilized by NODE O and NODE D are controlled by independent companies, it is unlikely that each company will utilize the same database software. Even if the companies did utilize the same database software, it is even more unlikely that each will have created a substantially identical record format structures because of the different needs of each company.

It is possible to address this problem by utilizing a single database which can be accessed by a variety of users. Such a solution has the disadvantage that all of the users must agree to utilize the same format and agree to a common set of rules regarding use and access of the database. It is difficult to achieve agreement among independent companies, or even different parts of the same company, because of their differing needs and objectives regarding database capabilities.

Difficulties would still exist even if independent companies utilized the same database software and record format. Although the common structure of the database records would enable records to be communicated over a conventional communication network and entered into another database, difficulties exist in maintaining duplicate records in different databases where the record can be updated by a user at any of the independent databases. Although it is possible to establish rules requiring that new and modified records be transmitted to others in the network, this places a substantial burden on each user in systems where it is desirable to maintain duplicate records at different databases so that information can be locally accessed.

Thus, there exists a need for a communication system which provides flexibility by allowing individual nodes to utilize different databases and which automatically

updates corresponding records at different databases with a minimum of burden on the users.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide an improved communication system which addresses the above needs.

In accordance with an embodiment of the present invention, a database interface and an associated communication program are installed at each node in a communication network. The database interface provides a translation between the record format used at a database and a common standardized record format. Thus, the database interface provides an inbound and outbound translation of fields within a record in order to provide compatibility with varying databases utilized at different nodes. The standardized format is used by the communication program to transmit record information to the other nodes in the system. The communication program provides a means for transmitting records and updates of records to the nodes in the network which need to maintain duplicate records. This communication capability is substantially automated to minimize user demands.

BRIEF DESCRIPTION OF THE DRAWING

FIG. 1 is a block diagram of a typical network which incorporates an embodiment of the present invention.

FIG. 2 is a block diagram of a typical node in the network of FIG. 1.

FIG. 3 is a diagram which illustrates the relationship between database interface software and communication software in accordance with the present invention.

FIG. 4 is a diagram which illustrates the mapping of a local database record into a standardized record format in accordance with the present invention.

FIG. 5 illustrates an embodiment of a communication directory and file structure in accordance with an embodiment of the present invention.

FIGS. 6 and 7 comprise a table that illustrates parameters utilized with each record in accordance with an embodiment of the present invention.

FIG. 8 is a flow diagram illustrating the transmission of a record from a node in accordance with an embodiment of the present invention.

FIG. 9 is a flow diagram illustrating receipt of a record from another node in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

An exemplary application of the present invention follows in order to appreciate the advantages of the present invention. Assume that a provider of telecommunications equipment is coupled via a communication system to a plurality of independent users of its equipment. A trouble ticket, i.e. a documented problem with a failure of user owned equipment, is entered by a user in an independent database controlled by the user. If the trouble ticket relates to equipment supplied by the manufacturer a copy of the trouble ticket is forwarded utilizing the database interface and communication program in accordance with an embodiment of the present invention to a corresponding field support organization of the equipment provider. If this organization is able to solve the problem, the duplicate ticket stored in a database of the support organization is modified indicating the solution or other relevant information relating to the

problem and transmitted to the originating user. If the organization is unable to solve the problem, it sends a copy of the ticket to a different organization within the company with greater technical expertise and sends an update of the same ticket back to the user so that the user is apprised of the status of the trouble ticket. Assuming the higher level organization solves or otherwise handles the problem, the higher level organization updates the trouble ticket and sends a copy of the updated trouble ticket back to the field support organization which in turn sends the update to the end user. Assuming the end user is satisfied with the solution to the problem indicated on the updated trouble ticket, the ticket is updated and closed by the user with copies of the update being sent to the field support organization and the higher level technical organization. Each organization dealing with the trouble ticket is apprised of all updates to the trouble ticket by any involved party. However, each of the parties (nodes in the communication system) are free to maintain an independent database so that each can independently maintain statistics and otherwise utilize the information according to the needs of each.

The illustrative embodiment of the present invention can be incorporated in a network utilized for problem solving as explained above. However, it will be apparent to those skilled in the art that the utility of the present invention is not limited to this application.

FIG. 1 illustrates a communication system in which a plurality of user nodes 10-20 are coupled to support nodes 22 and 24 which are in turn coupled to central support node 26. In the illustrative example, service providers at nodes 22 and 24 are in a direct support level and have sufficient expertise to solve many of the problems posed by the users which they support. Technical experts at node 26 provide corporate support to users for problems which the service providers at the direct support nodes cannot solve. The direct support nodes are also connected to each other so that problem solving expertise can be shared at the direct support level. A communication network 27 provides communication channels among the nodes.

This embodiment of the present invention is especially well suited for utilization in the illustrative system, especially when the equipment and/or software supplied to users is complex. For example, users may encounter difficulties in configuring the equipment to properly interface with other peripheral equipment or in customizing the operation of the equipment pursuant to a desired customer configuration. Help with such problems can normally be provided by the service providers at the support nodes. Another class of problems which users may encounter deal with failures or errors in the basic operation of the equipment. Some of these problems may be solved by personnel at the support nodes, but other problems will require the greater technical expertise available from the technical experts at the central support node.

In the illustrative example, a user documents a new problem by entering appropriate information describing the equipment, nature of the problem, and other parameters utilizing the user's database. The user preferred format for storing information in the local independent database is translated into a standardized record format by a database interface which is coupled to the local database. After the information is translated into the standardized record format, a communication program in accordance with the present invention interfaces

with the database interface and permits the standardized record to be transmitted to a direct support node. A database interface at the support node translates the standardized record into the appropriate format for the independent database maintained by the support node. The support node acknowledges receipt of the record to the user. Preferably, the direct support node generates an electronic mail (e-mail) message to an assigned support person, advising of the receipt of a problem ticket. After the direct support person considers the problem as documented by the user's record, the record will typically be updated by the assigned support person to contain additional information and may request additional information from the user relevant to the problem. The record at the support node is updated in the support node database and communicated to the originating user in standardized format. The database interface at the user node translates the received update from standardized format to the record format used at the user's local database. Thus, each node in the system which addresses a trouble ticket maintains an updated record regardless of the party making an update. Preferably, the person at the user node is notified by e-mail indicating that the record has been updated and thus, should be reviewed for possible additional action. Thus, the user and the assigned support person are each guaranteed of having access to the latest updated information concerning the problem while allowing each to maintain an independent database which can be tailored to meet the different needs of the user and the support provider.

FIG. 2 illustrates a typical node in a system in accordance with the present invention. The node includes a microprocessor (MPU) 28, and associated memory including read-only memory (ROM) 30, random access memory (RAM) 32, and alternate memory storage 34 such as a disk or tape drive. Conventional input and output devices associated with MPU 28 include a keyboard 36, a video display terminal 38, and a communication module 40 which enables communications between MPU 28 and a communication network 27. The node as shown in FIG. 2 can be used for any of the nodes as shown in FIG. 1. The communication network 27 may comprise either the public switched telephone network such as by modems or may comprise dedicated leased lines between the users and connected nodes.

FIG. 3 is a pictorial representation of the relationship among software in accordance with the present invention. Local database software 46 receives input from the user and provides output to the user of stored records. The database interface 44 translates the relevant fields in a record stored in the local database 46 into corresponding fields in a standardized record format. Records retrieved via interface 44 from the local database are likewise translated for transmission by communication program 48 to communication network 27 in the standardized record format. The standardized record format allows individual nodes to maintain different local database software programs and define fields within records to meet the needs of the associated users of the node. The local database may consist of conventionally available database software including flat file and relational types. In order to minimize the overall communication network requirements in order to handle the transmission of records, each of the nodes preferably utilizes a UNIX® operating system which inherently contains communication support.

FIG. 4 graphically illustrates how the database interface 44 interfaces with the local database 46. The standardized record format 50 is defined and stored in database interface 44. It is related to a local database record format 52 utilized by the local database 46. The standardized record format 50 includes fields A-N which define different information parameters which can be communicated within the network shown in FIG. 1. The local database record format 52 includes fields 1-99 which correspond to parameters determined by the operator of the node for possible storage in local database 46. The lines 54 which interconnect one field in format 50 with one field in format 52 identify the translation or mapping which will occur when transferring records 52 stored in local database 46 to other nodes in the system. Each of the fields in record format 50 are preferably supported by a corresponding field in record format 52. The standardized record format 50 which is transmitted between nodes allows the independent databases associated with each node to be able to share information with other nodes without requiring manual re-entry of the information at another database because of different record and field structures.

In a preferred embodiment of the present invention, one of the fields in standardized record format 50 identifies if the record contains information which should be interpreted by another node as a trouble ticket, i.e. typically a request for help by another node, and another field contains a description of the product or manufacturer of the product associated with the problem. Each node preferably maintains in memory an address map which identifies the support node to receive a copy of the record dependent upon the equipment manufacturer or product entered in the corresponding record field. Thus, a user as illustrated in FIG. 1 does not have to provide an address of the direct support node associated with the problem since this addressing will be accomplished automatically merely by the user initiating the trouble ticket. Since the database interface 44 is coupled to the communication program 48 and monitors changes made to records in the local database, a person entering data in local database 46 need only change the ownership field in a trouble ticket in order to initiate transmission of the record to the appropriate supporting node. It is contemplated that the user can be coupled by the communication network to a variety of different support nodes which support different products.

FIG. 5 illustrates an exemplary directory structure, preferably associated with each node, in accordance with the present invention, in order to facilitate communications. Node directory 56 is the highest level directory which interacts with communication network 27 to receive communications from the network which are addressed to the node and to transmit communications to another node in the network. Subdirectory BIN 58 contains the communication program 60 in accordance with the present invention. The communication program 60 is explained, especially with regard to the flow diagram of FIGS. 8 and 9. Subdirectories 62, 64, and 66 each have symmetrical file structures. Each of these subdirectories contains a structure in which one allocated subdirectory is maintained for each node in the system. Each record received by node directory 56 is placed under IN directory 62 and the subdirectory which corresponds to the node which transmitted the record. These files are periodically scanned by the database interface 44 to check for received files. Received

standardized records are translated utilizing the database interface 44 and then stored in the local database of the node. To transmit a record stored in the local database 46 of the node, the record is translated by the database interface 44 and a copy of the standardized record placed in the corresponding subdirectory of the destination node under OUT directory 64. The communication program 60 periodically scans the files in the OUT subdirectories and transmits the records via node directory 56 to the destination node in communication network 27. After the transmission of the outbound record, a duplicate copy of the record is placed in a corresponding subdirectory under HOLD directory 66; the transmitted record in the subdirectory under OUT directory 64 is deleted. If an acknowledgement is not received from the destination node within a predetermined time, an attempt to retransmit the record is repeated from the subdirectory under HOLD directory 66.

FIGS. 6 and 7 illustrate the status of parameters associated with a record relative to nodes 12, 22, and 26 of FIG. 1 at time intervals T0, T1, T2, T3, and T4. The state of the parameters shown in FIGS. 6 and 7 reflect states at the end of each interval. Time T0 corresponds to user node 12 creating a new record indicative of a problem. A unique ticket number which also identifies the originating node is assigned, i.e. UN 12-1487. User node 12 is also identified as the originator and owner of the trouble ticket; note the originator and owner rows in FIG. 6. The status of the trouble ticket is "open". In the illustrative example the record was created and entered in the local database of node 12 without causing the record to be transmitted as a trouble ticket to direct support node 22. This indicates that records may be opened and potentially closed by the originating user if the user determines a solution or otherwise decides that a trouble ticket is not to be generated. Thus the user's database can store records that are not trouble tickets.

Assuming that node 12 now wishes to communicate the stored record as a trouble ticket for direct support node 22, a person at node 12 will load the stored information from the local database 46 into database interface 44 which includes the parameters as shown. The owner is changed from node 12 to node 22. The ownership parameter identifies the node from which help is sought. The ownership of a given ticket may change from time to time depending upon the node from which help is sought. User node 12 always remains the originator and the ticket number is never changed during the communication of records from node to node. This allows each trouble ticket to be assigned a sequential number by each node. At user node 12, the record is placed in node 22 of the OUT directory 64 for transmission to node 22 by the communication program. Upon receiving the record at node 22, it is placed under IN directory 62 in the subdirectory corresponding to user node 12. It is then transferred to the local database 46 of node 22. Thus, following the successful completion of the transmission, the identical record will exist at nodes 12 and 22.

In this example, a service provider at node 22 is unable to handle the problem associated with ticket UN 12-1487 and makes a determination that the ticket should be escalated to node 26 for additional help from the technical experts in problem solving. At node 22, the record is retrieved from the local database into the database interface and the ownership changed from node 22 to node 26. Node 22 transmits the record at

time period T2 to node 26 in a similar manner explained as for the transmission from node 12 to node 22. In addition, node 22 transmits the updated record to node 12 so that all nodes in the network will maintain identical data for each ticket. Each node in the network maintains an address table which identifies the node from which it received each trouble ticket and thus, permits a daisy-chain addressing by which updated tickets are transmitted to the nodes which contain corresponding stored records. Node 26 receives the record and stores it in its local database. Preferably, upon receipt of each ticket at each node, a corresponding e-mail message is generated to the person at the node associated with handling such tickets. Thus, at the end of time interval T2, nodes 12, 22, and 26 will each contain an identical copy of the ticket in each's database.

After considering the substance of the problem indicated in the ticket, node 26 may make a determination of the cause of the problem and provide a solution. The documentation for the problem cause and the proposed solution will be included as updated information to the ticket during time interval T3 and the status changed by node 26 from "open" to "solved". Node 26 initiates an update transmission of the ticket to node 22 which in the daisy-chain manner retransmits the updated ticket back to the originating node 12. It will also be noted that at time interval T3 the owner is changed from node 26 to user node 12, indicating that further disposition of the problem ticket with the proposed solution rests with the user at node 12.

The user at node 12 considers the proposed solution and accepts the solution proposed. During time interval T4, the user at node 12 changes the status of the message from solved to closed and transmits the updated record to node 22 which in daisy-chain manner transmits the message to node 26. As illustrated in time T4, each of the nodes communicating relative to the illustrative trouble ticket will contain an identical record which indicates the status of the ticket is closed thereby ending further problem solving consideration of the item by direct support providers at node 22 and the technical experts at central support node 26. The record can be maintained in the local database of the respective nodes and utilized for other related problems should similar problems occur; the record can also be used for statistical data purposes.

FIG. 8 is a flow diagram illustrating steps in the transmission of a record to another node. Beginning at START 68, the record to be transmitted is loaded from the local database to the standard record format of the database interface 44 as indicated by step 70. It will be understood by those skilled in the art that the record may also be directly entered by the user into the database interface as well as being retrieved from storage in the local database. In step 72 the user requests transmission of the record to another node such as by changing the ownership of the record. The address of the destination node is determined in step 74. In a preferred embodiment of the present invention, the destination is determined from an originating node based upon a stored table which correlates the address of the node which services the equipment identified by the trouble ticket. Alternatively, the user at the originating node can directly enter the address of the node from which help is sought. Next, a copy of the trouble ticket is placed in the destination node (DN) subdirectory under the OUT directory. In accordance with the communication program, a predetermined time interval can be

utilized in which to check for multiple records to be transmitted to the same destination. Step 78 represents that the record is held waiting for an OUT transmission cycle to be started. At the end of the cycle, the record is sent to the destination node, as indicated in step 80 through communication network 27. In step 82 a copy of the transmitted record is placed in the destination node subdirectory under the HOLD directory. In decision step 84 a determination is made if an acknowledgement has been received from the destination node within a predetermined time. If NO, control transfers to the beginning of step 78 to initiate a retransmission of the record. Upon a YES determination by step 84, the record stored in the destination node subdirectory under HOLD is deleted since a retransmission will not be required. The transmission step terminates at END 88. Thus, in accordance with the transmission according to the present invention, a trouble ticket or record is transmitted to another node in standardized record format based upon a corresponding local database record format.

FIG. 9 is a flow diagram illustrating steps in accordance with the present invention for receiving a record transmitted from another node. Beginning at START 90, the general node directory 56 receives a record addressed for the subject node as indicated in step 92. In step 94 the received record is put under the IN directory in the subdirectory of the node which transmitted the record. In step 96 the record is copied from the subdirectory under the IN directory to the database interface 44 as a standardized record format. In step 95 an acknowledgement is sent to the node which transmitted the record acknowledging its successful receipt. In step 98 the database 46 is updated by the database interface 44 translating the standardized record format and storing the record in the local database 46. A determination is made in step 102 of whether to transmit the record to another node. This decision is based upon the destination contained with the record and the owner of the record. For example, if an originating node identifies the owner as a node other than an intermediate node which receives the record, the intermediate node will recognize the need to transmit the record to the designated owner. Upon a YES determination by step 102, control passes as indicated by transfer point "A" 104 back to step 74 as shown in FIG. 8. The transmission then continues as previously explained with regard to FIG. 8. A NO determination by step 102 indicates that further transmission is not required and concludes the receipt of the record at END 106.

In accordance with the present invention, a user friendly communication system is provided in which all nodes in the system requiring access to a record maintains the record in a local database. Subsequent updates of the record by any node are automatically distributed to the other nodes by utilizing a standardized record format. Thus, the present invention provides an enhanced communication system allowing independent database flexibility while still providing the relevant nodes in the network with up-to-date records.

Although an embodiment of the present invention has been described and shown in the drawings, the scope of the invention is defined by the claims which follow.

We claim:

1. A communication system comprising:
first and second nodes coupled to each other;
first and second databases associated with said first and second nodes respectively, each store records,

said first database storing said records in a first field format and said second database storing said records in a second field format which is different from said first field format;

means at said first and second nodes for transmitting records having standardized format to and receiving records having standardized format from the other of said first and second nodes, said standardized format being different from said first and second field formats;

means at said first node for automatically translating a received standardized format record into said first field format record for storage in said first database and for automatically translating a first field format record stored in said first database into standardized format record for transmission to said second node;

means for independently identifying an originator and an owner of each record, said translating and transmitting means acting to translate and transmit a first record contained in said first database associated with said first node to said second database associated with said second node based upon entry in said first record of an owner associated with said second database, wherein records in said first database that do not identify an owner are not transmitted to said second database;

said identifying means permitting the identified owner of a duplicate record stored at the second database and corresponding to said first record to be changed thereby changing responsibility for acting on the duplicate record and causing the duplicate record with changed ownership to be automatically transmitted to a database associated with the changed owner.

2. The system according to claim 1 further comprising a plurality of other nodes in addition to said first and second nodes, said other nodes comprising:

a database;

means for transmitting standardized format records and receiving standardized format records and means for automatically translating a received standardized format record into its field format record for storage in its database and for automatically translating a record stored in its database into a standardized format record for transmission to one of said first, second, and other nodes.

3. The system according to claim 1 further comprising means associated with said second node for automatically transmitting a copy of an updated first record based on the original record from said second database to the first database containing the corresponding original record, said first database storing said updated first record.

4. The system according to claim 1 further comprising means for labeling each record at its creation with a unique identification number which identifies all corresponding records in the communication system.

5. A first node for use in a communications system having a plurality of nodes, said first node comprising:

first database that stores records in a first field format;

means for transmitting records having standardized format to and receiving records having standardized format from one of said nodes, said standardized format being different from said first field format;

means for automatically translating a standardized format record into said first field format for storage

in said first database and for automatically translating a first field format record stored in said first database into a standardized format record for transmission to one of said nodes;

means for independently identifying an originator and an owner of each record, said translating and transmitting means acting to translate and transmit a first record contained in said first database associated with said first node to a second database associated with a second node based upon entry in said first record of an owner associated with said second database, wherein records in said first database that do not identify an owner are not transmitted to said second database;

said identifying means permitting the identified owner of a duplicate record stored at the second database and corresponding to said first record to be changed thereby changing responsibility for acting on the duplicate record and causing the duplicate record with changed ownership to be automatically transmitted to a database associated with the changed owner.

6. The node according to claim 5 further comprising means for automatically sending records stored in the first database that have been modified to other nodes which contain corresponding records thereby keeping corresponding records at other nodes updated.

7. The node according to claim 5 further comprising means for labeling each record at its creation with a unique identification number which always identifies all corresponding records in other nodes in the communication system.

8. A method for handling records in a communications system which includes at least first and second nodes having first and second databases, respectively, said first database storing records in a first field format and said second database storing records in a second field format which is different from said first field format, said method comprising the steps of:

transmitting records having a standardized format from said first node and receiving records having a standardized format at said first node from said second node, said standardized format being different from said first and second field formats;

automatically translating a standardized format record received by said first node into said first field format record for storage in said first database and for automatically translating a first field format record stored in said first database into a standardized format record for transmission to said second node;

independently identifying an originator and an owner of each record, said translating and transmitting steps translating and transmitting a first record contained in said first database associated with said first node to said second database associated with said second node based upon entry in said first record of an owner associated with said second database, wherein records in said first database that do not identify an owner are not transmitted to said second database;

said identifying step permitting the identified owner of a duplicate record stored at the second database and corresponding to said first record to be changed thereby changing responsibility for acting on the duplicate record and causing the duplicate record with changed ownership to be automati-

11

cally transmitted to a database associated with the changed owner.

9. The method according to claim 8 further comprising the step of automatically updating records stored in databases of other nodes upon a corresponding record being modified at one node in the communication system.

10. The method according to claim 8 further comprising the step of labeling each record at its creation with a unique identification number which identifies all corresponding records in other databases in the communication system.

11. The method according to claim 8 wherein said identifying step permits the owner of a corresponding set of records to be changed by any node which stores one of the corresponding records thereby allowing the responsibility for acting on the record to be changed.

12. A method for handling records by a node in a communications system having other nodes, said method comprising the steps of:

storing records using a first field format in a first database;

transmitting records having a standardized format to and receiving records having a standardized format from said other nodes, said standardized format being different from said first field format;

automatically translating a standardized format record received from one of said other nodes into a first field format record and storing said translated record in said first database;

translating a first field format record stored in said first database into standardized format record and transmitting said standardized record to at least one of said other nodes;

12

independently identifying an originator and an owner of each record, said translating and transmitting steps translating and transmitting a first record contained in said first database associated with said first node to said second database associated with said second node based upon entry in said first record of an owner associated with said second database, wherein records in said first database that do not identify an owner are not transmitted to said second database;

said identifying step permitting the identified owner of a duplicate record stored at the second database and corresponding to said first record to be changed thereby changing responsibility for acting on the duplicate record and causing the duplicate record with changed ownership to be automatically transmitted to a database associated with the changed owner.

13. The method according to claim 12 further comprising the step of automatically sending records stored in the first database that have been modified by a user to other nodes, which contain corresponding records thereby keeping corresponding records at other nodes updated.

14. The method according to claim 12 further comprising the step of labeling each record at its creation with a unique identification number which identifies all corresponding records in other nodes in the communication system.

15. The method according to claim 12 wherein said identifying step permits the owner of a corresponding set of records to be changed by a node which stores one of the corresponding records thereby allowing the responsibility for acting on the record to be changed.

* * * * *